



**7450 ETHERNET SERVICE SWITCH
7750 SERVICE ROUTER
7950 EXTENSIBLE ROUTING SYSTEM
VIRTUALIZED SERVICE ROUTER**

**LAYER 3 SERVICES GUIDE: IES AND VPRN
SERVICES
RELEASE 16.0.R5**

3HE 14131 AAAC TQZZA 01

Issue: 01

December 2018

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

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

© 2018 Nokia.

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

Table of Contents

1	Getting Started.....	11
1.1	About This Guide.....	11
1.2	Layer 3 Services Configuration Process	13
2	Internet Enhanced Service.....	15
2.1	IES Service Overview.....	15
2.2	IES Features.....	17
2.2.1	IP Interfaces	17
2.2.1.1	QoS Policy Propagation Using BGP (QPPB)	17
2.2.1.2	QPPB.....	19
2.2.1.3	QPPB and GRT Lookup	24
2.2.1.4	Object Grouping and State Monitoring	27
2.2.2	Subscriber Interfaces.....	29
2.2.2.1	IPv6 Enhanced Subscriber Management (ESM).....	29
2.2.2.2	RADIUS Accounting	30
2.2.3	SAPs.....	32
2.2.3.1	Encapsulations	32
2.2.3.2	ATM SAP Encapsulations for IES	33
2.2.3.3	Pseudowire SAPs.....	33
2.2.3.4	Encapsulation	34
2.2.3.5	Pseudowire SAP Configuration	35
2.2.3.6	QoS for Pseudowire Ports and Pseudowire SAPs	36
2.2.3.7	Shaping and Bandwidth Control	37
2.2.3.8	Lag Considerations.....	39
2.2.3.9	Last Mile Packet Size Adjustment	39
2.2.3.10	Redundancy with Pseudowire SAPs	40
2.2.3.11	Operational Group Support for PW Ports	42
2.2.4	Routing Protocols	43
2.2.4.1	CPE Connectivity Check	43
2.2.5	QoS Policies	44
2.2.6	Filter Policies	44
2.2.7	MPLS Entropy Label and Hash Label	44
2.2.8	Spoke SDPs	44
2.2.9	SRRP.....	46
2.2.9.1	SRRP Messaging	51
2.2.9.2	SRRP and Multi-Chassis Synchronization	53
2.2.9.3	SRRP Instance	53
2.2.9.4	Subscriber Subnet Owned IP Address Connectivity	57
2.2.9.5	Subscriber Subnet SRRP Gateway IP Address Connectivity.....	57
2.2.9.6	Receive SRRP Advertisement SAP and Anti-Spoof.....	57
2.2.9.7	BFD with SRRP/VRRP	58
2.3	Configuring an IES Service with CLI	59
2.3.1	Basic Configuration	59
2.3.2	Common Configuration Tasks.....	59
2.3.3	Configuring IES Components.....	60

2.3.3.1	Configuring an IES Service	60
2.3.3.2	Configuring IES Subscriber Interface Parameters.....	61
2.3.3.3	Configuring IES Interface Parameters.....	61
2.3.3.4	Configuring Spoke-SDP Parameters.....	62
2.3.3.5	Configuring SAP Parameters	62
2.3.3.6	Configuring IES SAP ATM Parameters	63
2.3.3.7	Configuring VRRP	65
2.3.3.8	Configuring IPSec Parameters	65
2.3.3.9	IGMP Host Tracking	66
2.4	Service Management Tasks	67
2.4.1	Modifying IES Service Parameters.....	67
2.4.2	Deleting a Spoke-SDP.....	68
2.4.3	Deleting an IES Service.....	68
2.4.4	Disabling an IES Service	69
2.4.5	Re-Enabling an IES Service	69
2.5	IES Services Command Reference.....	71
2.5.1	Command Hierarchies.....	71
2.5.1.1	Global Commands.....	71
2.5.1.2	AARP Interface Commands	71
2.5.1.3	IES Service Interface Commands	72
2.5.1.4	Routed VPLS Commands	77
2.5.1.5	Redundant Interface Commands.....	78
2.5.1.6	Interface SAP Commands	78
2.5.1.7	VRRP Commands	84
2.5.1.8	Spoke SDP Commands.....	85
2.5.1.9	Subscriber Interface Commands	88
2.5.2	Command Descriptions	96
2.5.2.1	Generic Commands.....	96
2.5.2.2	IES Global Commands	99
2.5.2.3	AARP Interface Commands	103
2.5.2.4	IES Interface Commands	106
2.5.2.5	Redundant Interface Commands.....	162
2.5.2.6	IES Subscriber Interface Commands	163
2.5.2.7	IES Interface DHCP Commands	169
2.5.2.8	PPPoE Commands.....	180
2.5.2.9	IES Interface ICMP Commands	182
2.5.2.10	IES Interface IPv6 Commands	189
2.5.2.11	IES Spoke SDP Commands.....	209
2.5.2.12	IES SAP Commands	221
2.5.2.13	SAP Subscriber Management Commands.....	237
2.5.2.14	ETH-CFM Service Commands	244
2.5.2.15	IES Filter and QoS Policy Commands.....	256
2.5.2.16	ATM Commands.....	289
2.5.2.17	IES Interface VRRP Commands	293
2.5.2.18	IPSec Gateway Commands	303
2.6	IES Show, Clear, and Debug Command Reference	309
2.6.1	Command Hierarchies.....	309
2.6.1.1	Show Commands	309
2.6.1.2	Clear Commands.....	311

2.6.1.3	Debug Commands.....	311
2.6.1.4	Monitor Commands	312
2.6.2	Command Descriptions	312
2.6.2.1	IES Show Commands	312
2.6.2.2	IES Clear Commands.....	378
2.6.2.3	IES Debug Commands.....	386
2.6.2.4	IES Monitor Commands	388
3	Virtual Private Routed Network Service	389
3.1	VPRN Service Overview.....	389
3.1.1	Routing Prerequisites	390
3.1.2	Core MP-BGP Support.....	390
3.1.3	Route Distinguishers	391
3.1.3.1	eiBGP Load Balancing	392
3.1.4	Route Reflector.....	393
3.1.5	CE to PE Route Exchange	393
3.1.5.1	Route Redistribution.....	394
3.1.5.2	CPE Connectivity Check	394
3.1.6	Constrained Route Distribution (RT Constraint).....	396
3.1.6.1	Constrained VPN Route Distribution Based on Route Targets	396
3.1.6.2	Configuring the Route Target Address Family	396
3.1.6.3	Originating RT Constraint Routes.....	397
3.1.6.4	Receiving and Re-Advertising RT Constraint Routes.....	397
3.1.6.5	Using RT Constraint Routes.....	398
3.1.7	BGP Fast Reroute in a VPRN	401
3.1.7.1	BGP Fast Reroute in a VPRN Configuration	402
3.1.8	BGP Best-External in a VPRN Context	402
3.2	VPRN Features	404
3.2.1	IP Interfaces	404
3.2.1.1	QoS Policy Propagation Using BGP (QPPB)	404
3.2.1.2	QPPB Applications	405
3.2.1.3	Inter-AS Coordination of QoS Policies	405
3.2.1.4	Traffic Differentiation Based on Route Characteristics	406
3.2.1.5	QPPB.....	406
3.2.1.6	Associating an FC and Priority with a Route	407
3.2.1.7	Displaying QoS Information Associated with Routes	408
3.2.1.8	Enabling QPPB on an IP interface	409
3.2.1.9	QPPB When Next-Hops are Resolved by QPPB Routes.....	410
3.2.1.10	QPPB and Multiple Paths to a Destination	410
3.2.1.11	QPPB and Policy-Based Routing	411
3.2.1.12	QPPB and GRT Lookup	411
3.2.1.13	QPPB Interaction with SAP Ingress QoS Policy.....	411
3.2.1.14	Object Grouping and State Monitoring	414
3.2.1.15	VPRN IP Interface Applicability	414
3.2.2	Subscriber Interfaces.....	416
3.2.3	SAPs.....	416
3.2.3.1	Encapsulations	416
3.2.3.2	ATM SAP Encapsulations for VPRN Services	417
3.2.3.3	Pseudowire SAPs.....	417

3.2.4	QoS Policies	417
3.2.5	Filter Policies	418
3.2.6	DSCP Marking.....	418
3.2.6.1	Default DSCP Mapping Table	419
3.2.7	Configuration of TTL Propagation for VPRN Routes	420
3.2.8	CE to PE Routing Protocols	422
3.2.8.1	PE to PE Tunneling Mechanisms	422
3.2.8.2	Per VRF Route Limiting.....	422
3.2.9	Spoke SDPs	423
3.2.9.1	T-LDP Status Signaling for Spoke-SDPs Terminating on IES/VPRN.....	424
3.2.9.2	Spoke SDP Redundancy into IES/VPRN	424
3.2.10	IP-VPNs.....	425
3.2.10.1	Using OSPF in IP-VPNs	425
3.2.11	IPCP Subnet Negotiation.....	426
3.2.12	Cflowd for IP-VPNs.....	427
3.2.13	Inter-AS VPRNs.....	427
3.2.14	VPRN Label Security at Inter-AS Boundary	430
3.2.14.1	Feature Configuration.....	431
3.2.14.2	CPM Behavior	432
3.2.14.3	Data Path Forwarding Behavior	433
3.2.15	Carrier Supporting Carrier (CSC)	434
3.2.15.1	Terminology.....	435
3.2.15.2	CSC Connectivity Models.....	436
3.2.15.3	CSC-PE Configuration and Operation.....	436
3.2.15.4	CSC Interface	436
3.2.15.5	QoS	438
3.2.15.6	MPLS.....	439
3.2.15.7	CSC VPRN Service Configuration.....	440
3.2.16	Node Management Using VPRN	440
3.2.16.1	VPRN Management.....	441
3.2.16.2	AAA Management	441
3.2.17	Traffic Leaking to GRT	442
3.2.18	Traffic Leaking from VPRN to GRT for IPv6.....	444
3.2.19	RIP Metric Propagation in VPRNs.....	444
3.2.20	NTP Within a VPRN Service	445
3.2.21	PTP Within a VPRN Service.....	445
3.2.22	VPN Route Label Allocation	445
3.2.22.1	Configuring the Service Label Mode	447
3.2.22.2	Restrictions and Usage Notes	447
3.2.23	VPRN Support for BGP Flowspec.....	448
3.2.24	MPLS Entropy Label and Hash Label	448
3.2.25	LSP Tagging for BGP Next-Hops or Prefixes and BGP-LU	449
3.2.26	Route Leaking from the Global Route Table to VPRN Instances	449
3.3	QoS on Ingress Bindings.....	450
3.4	Multicast in IP-VPN Applications	452
3.4.1	Use of Data MDTs	453
3.4.2	Multicast Protocols Supported in the Provider Network	454
3.4.3	MVPN Membership Auto-discovery using BGP	455

3.4.4	PE-PE Transmission of C-Multicast Routing using BGP	457
3.4.5	VRF Route Import Extended Community	457
3.4.6	Provider Tunnel Support.....	458
3.4.6.1	Point-to-Multipoint Inclusive (I-PMSI) and Selective (S-PMSI) Provider Multicast Service Interface	458
3.4.6.2	P2MP RSVP-TE I-PMSI and S-PMSI	459
3.4.6.3	P2MP LDP I-PMSI and S-PMSI	459
3.4.6.4	Wildcard (C-*, C-*) P2MP LSP S-PMSI	460
3.4.6.5	P2MP LSP S-PMSI.....	463
3.4.6.6	Dynamic Multicast Signaling over P2MP LDP in VRF	464
3.4.6.7	MVPN Sender-only/Receiver-only	465
3.4.6.8	S-PMSI Trigger Thresholds	467
3.4.6.9	Migration from Existing Rosen Implementation	468
3.4.6.10	Policy-based S-PMSI.....	468
3.4.6.11	Policy-based Data MDT.....	472
3.4.7	MVPN (NG-MVPN) Upstream Multicast Hop Fast Failover.....	473
3.4.8	Multicast VPN Extranet.....	474
3.4.8.1	Multicast Extranet for Rosen MVPN for PIM SSM.....	474
3.4.8.2	Multicast Extranet for NG-MVPN for PIM SSM	475
3.4.8.3	Multicast Extranet with Per-group Mapping for PIM SSM	476
3.4.8.4	Multicast GRT-source/VRF-receiver Extranet with Per Group Mapping for PIM SSM	478
3.4.8.5	Multicast Extranet with Per-group Mapping for PIM ASM	479
3.4.9	Non-Congruent Unicast and Multicast Topologies for Multicast VPN	481
3.4.10	Automatic Discovery of Group-to-RP Mappings (Auto-RP).....	482
3.4.11	IPv6 MVPN Support	483
3.4.12	Multicast Core Diversity for Rosen MDT_SAFI MVPNs	485
3.4.13	NG-MVPN Core Diversity	487
3.4.13.1	NG-MVPN to Loopback Interface	488
3.4.13.2	NG-MVPN Core Diversity	489
3.4.13.3	Configuration Example	491
3.4.14	NG-MVPN Multicast Source Geo-Redundancy.....	493
3.4.15	Multicast Core Diversity for Rosen MDT SAFI MVPNs	496
3.4.16	Inter-AS MVPN	498
3.4.16.1	BGP Connector Attribute	499
3.4.16.2	PIM RPF Vector.....	499
3.4.16.3	Inter-AS MVPN Option B	500
3.4.16.4	Inter-AS MVPN Option C	501
3.4.16.5	NG-MVPN Non-segmented Inter-AS Solution.....	503
3.4.17	Weighted ECMP and ECMP for VPRN IPv4 and IPv6 over MPLS LSPs.....	513
3.5	FIB Prioritization	514
3.6	Configuring a VPRN Service with CLI	515
3.6.1	Basic Configuration	515
3.6.2	Common Configuration Tasks	516
3.6.3	Configuring VPRN Components.....	517
3.6.3.1	Creating a VPRN Service	517
3.6.3.2	Configuring Global VPRN Parameters	518

3.6.3.3	Configuring VPRN Log Parameters.....	518
3.6.3.4	Configuring VPRN Protocols - PIM.....	520
3.6.4	Configuring IPsec Parameters	531
3.7	Service Management Tasks	532
3.7.1	Modifying VPRN Service Parameters.....	532
3.7.2	Deleting a VPRN Service	533
3.7.3	Disabling a VPRN Service	533
3.7.4	Re-enabling a VPRN Service	534
3.8	VPRN Service Configuration Commands.....	535
3.8.1	Command Hierarchies.....	535
3.8.1.1	VPRN Service Configuration Commands.....	536
3.8.1.2	L2TP Commands.....	540
3.8.1.3	DHCP Commands	543
3.8.1.4	GSMP Commands.....	546
3.8.1.5	IGMP Commands	547
3.8.1.6	IPsec Configuration Commands.....	548
3.8.1.7	Log Commands	551
3.8.1.8	Remote Server Commands	553
3.8.1.9	VPRN Management Commands	554
3.8.1.10	Multicast VPN Commands.....	554
3.8.1.11	Redundant Interface Commands.....	556
3.8.1.12	Router Advertisement Commands	557
3.8.1.13	NTP Commands	558
3.8.1.14	NAT Commands	558
3.8.1.15	Subscriber Interface Commands	559
3.8.1.16	Interface Commands	567
3.8.1.17	Network Interface Commands	571
3.8.1.18	Interface Spoke SDP Commands.....	573
3.8.1.19	Interface VRRP Commands	575
3.8.1.20	Interface SAP Commands	576
3.8.1.21	Interface SAP Tunnel Commands	581
3.8.1.22	Routed VPLS Commands	585
3.8.1.23	Oper Group Commands	585
3.8.1.24	Network Ingress Commands	586
3.8.1.25	BGP Configuration Commands	586
3.8.1.26	BGP Group Configuration Commands	589
3.8.1.27	BGP Group Neighbor Configuration Commands.....	591
3.8.1.28	IS-IS Configuration Commands	593
3.8.1.29	OSPF Configuration Commands	597
3.8.1.30	PIM Configuration Commands	603
3.8.1.31	MSDP Configuration Commands	606
3.8.1.32	MLD Configuration Commands	607
3.8.1.33	RIP Configuration Commands.....	608
3.8.1.34	RADIUS Commands.....	610
3.8.1.35	Web Portal Protocol Configuration Commands.....	611
3.8.1.36	AARP Interface Commands	612
3.8.2	Command Descriptions	613
3.8.2.1	Generic Commands.....	613
3.8.2.2	Global Commands.....	616

3.8.2.3	Router L2TP Commands.....	641
3.8.2.4	Router DHCP Configuration Commands.....	661
3.8.2.5	IGMP Commands.....	683
3.8.2.6	IPSec Configuration Commands.....	730
3.8.2.7	Log Commands.....	739
3.8.2.8	Remote Server Commands.....	757
3.8.2.9	VPRN Management Commands.....	766
3.8.2.10	Multicast VPN Commands.....	768
3.8.2.11	Redundant Interface Commands.....	789
3.8.2.12	Router Advertisement Commands.....	789
3.8.2.13	Network Time Protocol Commands.....	796
3.8.2.14	NAT Commands.....	799
3.8.2.15	Subscriber Interface Commands.....	808
3.8.2.16	Interface Commands.....	812
3.8.2.17	Network Interface Commands.....	859
3.8.2.18	Interface DHCP Commands.....	859
3.8.2.19	Interface ICMP Commands.....	876
3.8.2.20	Interface SAP ATM Commands.....	882
3.8.2.21	Interface Anti-Spoofing Commands.....	885
3.8.2.22	Interface SAP Filter and QoS Policy Commands.....	898
3.8.2.23	Interface VRRP Commands.....	939
3.8.2.24	Interface SAP Commands.....	949
3.8.2.25	Routed VPLS Commands.....	956
3.8.2.26	Network Ingress Commands.....	962
3.8.2.27	SAP Subscriber Management Commands.....	964
3.8.2.28	BGP Commands.....	972
3.8.2.29	ETH-CFM Service Commands.....	1019
3.8.2.30	IS-IS Commands.....	1031
3.8.2.31	OSPF Commands.....	1078
3.8.2.32	PIM Commands.....	1122
3.8.2.33	PPPoE Commands.....	1146
3.8.2.34	MSDP Configuration Commands.....	1148
3.8.2.35	MLD Configuration Commands.....	1157
3.8.2.36	RIP Commands.....	1164
3.8.2.37	SDP Commands.....	1175
3.8.2.38	RADIUS Proxy Commands.....	1188
3.8.2.39	AARP Interface Commands.....	1198
3.9	VPRN Show, Clear, and Debug Command Reference.....	1203
3.9.1	Command Hierarchies.....	1203
3.9.1.1	Show Commands.....	1203
3.9.1.2	Clear Commands.....	1206
3.9.1.3	Debug Commands.....	1207
3.9.2	Command Descriptions.....	1207
3.9.2.1	VPRN Show Commands.....	1207
3.9.2.2	VPRN Clear Commands.....	1338
3.9.2.3	VPRN Debug Commands.....	1348
3.10	Tools Command Reference.....	1355
3.10.1	Command Hierarchies.....	1355
3.10.1.1	Tools Commands.....	1355

3.10.2	Command Descriptions	1355
3.10.2.1	Tools Commands	1355
4	Standards and Protocol Support	1359

1 Getting Started

1.1 About This Guide

This guide describes Layer 3 service functionality provided by Nokia's family of routers and presents examples to configure and implement various protocols and services.

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

The topics and commands described in this document apply to the:

- 7450 ESS
- 7750 SR
- 7950 XRS
- VSR

[Table 1](#) lists the available chassis types for each SR OS router.

Table 1 Supported SR OS Router Chassis Types

7450 ESS	7750 SR	7950 XRS
• 7450 ESS-7/12 running in standard mode (not mixed-mode)	• 7450 ESS-7/12 running in mixed-mode (not standard mode) • 7750 SR-a4/a8 • 7750 SR-c4/c12 • 7750 SR-1e/2e/3e • 7750 SR-7/12 • 7750 SR-12e	• 7950 XRS-16c • 7950 XRS-20/40

For a list of unsupported features by platform and chassis, refer to the *SR OS 16.0.Rx Software Release Notes*, part number 3HE 14220 000x TQZZA or the *VSR Release Notes*, part number 3HE 14204 000x TQZZA.

Command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.



Note: This guide generically covers Release 16.0.Rx content and may contain some content that will be released in later maintenance loads. Please refer to the *SR OS 16.0.Rx* Software Release Notes, part number 3HE 14220 000x TQZZA or the *VSR Release Notes*, part number 3HE 14204 000x, for information about features supported in each load of the Release 16.0.Rx software.

1.2 Layer 3 Services Configuration Process

[Table 2](#) lists tasks related to the configuration and implementation of Layer 3 Services.

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

Table 2 Configuration Process

Area	Task	Section
Internet Enhanced Service (IES)	Configure an IES service	Configuring an IES Service with CLI
	Configure IES components	Configuring IES Components
	Service management	Service Management Tasks
Virtual Private Routed Network (VPRN) Service	Configure a VPRN service	Configuring a VPRN Service with CLI
	Configure VPRN components	Configuring VPRN Components
	Service management	Service Management Tasks

2 Internet Enhanced Service


2.1 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 customer-facing IP interfaces to participate in the same routing instance used for service network core routing connectivity. IES services require that the IP addressing scheme used by the subscriber be unique between other provider addressing schemes and potentially the entire Internet.

While IES is part of the routing domain, the usable IP address space may be limited. This allows a portion of the service provider address space to be reserved for service IP provisioning, and be administered by a separate but subordinate address authority.

IP interfaces defined within the context of an IES service must have a SAP associated as the uplink access point to the subscriber network. Multiple IES services are created to segregate subscriber-owned IP interfaces.

Figure 1 Internet Enhanced Service



OSSG023

The IES service provides Internet connectivity. Other features include:

- multiple IES services are created to separate customer-owned 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.

These features apply to the 7750 SR and 7450 ESS.

Refer to the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* for information about how subscriber group-interfaces function in the Routed Central Office model.

2.2 IES Features

This section describes the 7450 ESS and 7750 SR service features and any special capabilities or considerations as they relate to IES services.

2.2.1 IP Interfaces

IES customer IP interfaces can be configured with most of the same options found on the core IP interfaces. The advanced configuration options supported are:

- QoS Policy Propagation Using BGP (QPPB)
- VRRP - for IES services with more than one IP interface
- Cflowd
- Secondary IP addresses
- ICMP Options

Configuration options found on core IP interfaces not supported on IES IP interfaces are:

- MPLS forwarding
- NTP broadcast receipt

2.2.1.1 QoS Policy Propagation Using BGP (QPPB)

This section discusses QPPB as it applies to VPRN, IES, and router interfaces. Refer to the [Internet Enhanced Service](#) section and the “IP Router Configuration” section in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide*.

QoS policy propagation using BGP (QPPB) is a feature that allows a route to be installed in the routing table with a forwarding-class and priority so that packets matching the route can receive the associated QoS. The forwarding-class and priority associated with a BGP route are set using BGP import route policies. In the industry, this feature is called QPPB, and even though the feature name refers to BGP specifically. On SR OS, QPPB is supported for BGP (IPv4, IPv6, VPN-IPv4, VPN-IPv6), RIP and static routes.

While SAP ingress and network QoS policies can achieve the same end result as QPPB, the effort involved in creating the QoS policies, keeping them up-to-date, and applying them across many nodes is much greater than with QPPB. This is due to assigning a packet, arriving on a particular IP interface, to a specific forwarding-class and priority/profile, based on the source IP address or destination IP address of the packet. In a typical application of QPPB, a BGP route is advertised with a BGP community attribute that conveys a particular QoS. Routers that receive the advertisement accept the route into their routing table and set the forwarding-class and priority of the route from the community attribute.

2.2.1.1.1 QPPB Applications

There are two typical applications of QPPB:

1. coordination of QoS policies between different administrative domains
2. traffic differentiation within a single domain, based on route characteristics

2.2.1.1.2 Inter-AS Coordination of QoS Policies

The operator of an administrative domain A can use QPPB to signal to a peer administrative domain B that traffic sent to certain prefixes advertised by domain A should receive a particular QoS treatment in domain B. More specifically, an ASBR of domain A can advertise a prefix XYZ to domain B and include a BGP community attribute with the route. The community value implies a particular QoS treatment, as agreed by the two domains (in their peering agreement or service level agreement, for example). When the ASBR and other routers in domain B accept and install the route for XYZ into their routing table, they apply a QoS policy on selected interfaces that classifies traffic towards network XYZ into the QoS class implied by the BGP community value.

QPPB may also be used to request that traffic sourced from certain networks receive appropriate QoS handling in downstream nodes that may span different administrative domains. This can be achieved by advertising the source prefix with a BGP community, as discussed above. However, in this case other approaches are equally valid, such as marking the DSCP or other CoS fields based on source IP address so that downstream domains can take action based on a common understanding of the QoS treatment implied by different DSCP values.


In the above examples, coordination of QoS policies using QPPB could be between a business customer and its IP VPN service provider, or between one service provider and another.

2.2.1.1.3 Traffic Differentiation Based on Route Characteristics

There may be times when a network operator wants to provide differentiated service to certain traffic flows within its network, and these traffic flows can be identified with known routes. For example, the operator of an ISP network may want to give priority to traffic originating in a particular ASN (the ASN of a content provider offering over-the-top services to the ISP's customers), following a certain AS_PATH, or destined for a particular next-hop (remaining on-net vs. off-net).

Figure 2 shows an example of an ISP that has an agreement with the content provider managing AS300 to provide traffic sourced and terminating within AS300 with differentiated service appropriate to the content being transported. In this example, we presume that ASBR1 and ASBR2 mark the DSCP of packets terminating and sourced, respectively, in AS300 so that other nodes within the ISP's network do not need to rely on QPPB to determine the correct forwarding-class to use for the traffic. The DSCP or other CoS markings could be left unchanged in the ISP's network and QPPB used on every node.

Figure 2 Use of QPPB to Differentiate Traffic in an ISP Network



2.2.1.2 QPPB

There are two main aspects of the QPPB feature:

- the ability to associate a forwarding-class and priority with certain routes in the routing table; and
- The ability to classify an IP packet arriving on a particular IP interface to the forwarding-class and priority associated with the route that best matches the packet.

2.2.1.2.1 Associating an FC and Priority with a Route

This feature uses a command in the route-policy hierarchy to set the forwarding class and optionally the priority associated with routes accepted by a route-policy entry. The command has the following structure:

fc *fc-name* [priority {low | high}]

The use of this command is illustrated by the following example:

```
config>router>policy-options
begin
community gold members 300:100
policy-statement qppb_policy
entry 10
from
protocol bgp
community gold
exit
action accept
fc hl priority high
exit
exit
exit
commit
```

The **fc** command is supported with all existing from and to match conditions in a route policy entry and with any action other than reject, it is supported with next-entry, next-policy and accept actions. If a next-entry or next-policy action results in multiple matching entries, then the last entry with a QPPB action determines the forwarding class and priority.

A route policy that includes the **fc** command in one or more entries can be used in any import or export policy but the **fc** command has no effect except in the following types of policies:

- VRF import policies:
 - config>service>vprn>vrf-import
- BGP import policies:
 - config>router>bgp>import
 - config>router>bgp>group>import
 - config>router>bgp>group>neighbor>import
 - config>service>vprn>bgp>import
 - config>service>vprn>bgp>group>import
 - config>service>vprn>bgp>group>neighbor>import
- RIP import policies:

- config>router>rip>import
- config>router>rip>group>import
- config>router>rip>group>neighbor>import
- config>service>vprn>rip>import
- config>service>vprn>rip>group>import
- config>service>vprn>rip>group>neighbor>import

As evident from above, QPPB route policies support routes learned from RIP and BGP neighbors of a VPRN as well as for routes learned from RIP and BGP neighbors of the base/global routing instance.

QPPB is supported for BGP routes belonging to any of the address families listed below:

- IPv4 (AFI=1, SAFI=1)
- IPv6 (AFI=2, SAFI=1)
- VPN-IPv4 (AFI=1, SAFI=128)
- VPN-IPv6 (AFI=2, SAFI=128)

A VPN-IP route may match both a VRF import policy entry and a BGP import policy entry (if vpn-apply-import is configured in the base router BGP instance). In this case the VRF import policy is applied first and then the BGP import policy, so the QPPB QoS is based on the BGP import policy entry.

This feature also introduces the ability to associate a forwarding-class and optionally priority with IPv4 and IPv6 static routes. This is achieved by specifying the forwarding-class within the static-route-entry next-hop or indirect context.

Priority is optional when specifying the forwarding class of a static route, but once configured it can only be deleted and returned to unspecified by deleting the entire static route.

2.2.1.2.2 Displaying QoS Information Associated with Routes

The following commands are enhanced to show the forwarding-class and priority associated with the displayed routes:

- **show router route-table**
- **show router fib**
- **show router bgp routes**
- **show router rip database**

- **show router static-route**

This feature uses a **qos** keyword to the **show>router>route-table** command. When this option is specified the output includes an additional line per route entry that displays the forwarding class and priority of the route. If a route has no fc and priority information then the third line is blank. The following CLI shows an example:

**show router route-table [family] [ip-prefix[/prefix-length]] [longer | exact]
[protocol protocol-name] qos**

An example output of this command is shown below:

```
A:Dut-A# show router route-table 10.1.5.0/24 qos
=====
Route Table (Router: Base)
=====
Dest Prefix                                Type    Proto    Age          Pref
      Next Hop[Interface Name]              Metric
      QoS
-----
10.1.5.0/24                                Remote   BGP       15h32m52s    0
      PE1_to_PE2
      h1, high
-----
No. of Routes: 1
=====
A:Dut-A#
```

2.2.1.2.3 Enabling QPPB on an IP Interface

To enable QoS classification of ingress IP packets on an interface based on the QoS information associated with the routes that best match the packets the **qos-route-lookup** command is necessary in the configuration of the IP interface. The **qos-route-lookup** command has parameters to indicate whether the QoS result is based on lookup of the source or destination IP address in every packet. There are separate **qos-route-lookup** commands for the IPv4 and IPv6 packets on an interface, which allows QPPB to be enabled for IPv4 only, IPv6 only, or both IPv4 and IPv6. The current QPPB based on a source IP address is not supported for IPv6 packets nor is it supported for ingress subscriber management traffic on a group interface.

The **qos-route-lookup** command is supported on the following types of IP interfaces:

- base router network interfaces (config>router>interface)
- VPRN SAP and spoke SDP interfaces (config>service>vprn>interface)
- VPRN group-interfaces (config>service>vprn>sub-if>grp-if)
- IES SAP and spoke SDP interfaces (config>service>ies>interface)

- IES group-interfaces (config>service>ies>sub-if>grp-if)

When the **qos-route-lookup** command with the destination parameter is applied to an IP interface and the destination address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the destination address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network qos policy.

Similarly, when the **qos-route-lookup** command with the source parameter is applied to an IP interface and the source address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the source address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network qos policy.

Currently, QPPB is not supported for ingress MPLS traffic on network interfaces or on CsC PE'-CE' interfaces (config>service>vpn>nw-if).



Note: QPPB based on a source IP address is not supported for ingress subscriber management traffic on a group interface.

2.2.1.2.4 QPPB When Next-Hops are Resolved by QPPB Routes

In some circumstances (IP VPN inter-AS model C, Carrier Supporting Carrier, indirect static routes, etc.) an IPv4 or IPv6 packet may arrive on a QPPB-enabled interface and match a route A1 whose next-hop N1 is resolved by a route A2 with next-hop N2 and perhaps N2 is resolved by a route A3 with next-hop N3, etc. The QPPB result is based only on the forwarding-class and priority of route A1. If A1 does not have a forwarding-class and priority association then the QoS classification is not based on QPPB, even if routes A2, A3, etc. have forwarding-class and priority associations.

2.2.1.2.5 QPPB and Multiple Paths to a Destination

When ECMP is enabled some routes may have multiple equal-cost next-hops in the forwarding table. When an IP packet matches such a route the next-hop selection is typically based on a hash algorithm that tries to load balance traffic across all the next-hops while keeping all packets of a given flow on the same path. The QPPB configuration model described in [Associating an FC and Priority with a Route](#) allows different QoS information to be associated with the different ECMP next-hops of a route. The forwarding-class and priority of a packet matching an ECMP route is based on the particular next-hop used to forward the packet.

When BGP fast reroute [1] is enabled some BGP routes may have a backup next-hop in the forwarding table in addition to the one or more primary next-hops representing the equal-cost best paths allowed by the ECMP/multipath configuration. When an IP packet matches such a route a reachable primary next-hop is selected (based on the hash result) but if all the primary next-hops are unreachable then the backup next-hop is used. The QPPB configuration model described in [Associating an FC and Priority with a Route](#) allows the forwarding-class and priority associated with the backup path to be different from the QoS characteristics of the equal-cost best paths. The forwarding class and priority of a packet forwarded on the backup path is based on the **fc** and priority of the backup route.

2.2.1.2.6 QPPB and Policy-Based Routing

When an IPv4 or IPv6 packet with destination address X arrives on an interface with both QPPB and policy-based-routing enabled:

- There is no QPPB classification if the IP filter action redirects the packet to a directly connected interface, even if X is matched by a route with a forwarding-class and priority
- QPPB classification is based on the forwarding-class and priority of the route matching IP address Y if the IP filter action redirects the packet to the indirect next-hop IP address Y, even if X is matched by a route with a forwarding-class and priority

2.2.1.3 QPPB and GRT Lookup

Source-address based QPPB is not supported on any SAP or spoke SDP interface of a VPRN configured with the **grt-lookup** command.

2.2.1.3.1 QPPB Interaction with SAP Ingress QoS Policy

When QPPB is enabled on a SAP IP interface the forwarding class of a packet may change from **fc1**, the original **fc** determined by the SAP ingress QoS policy to **fc2**, the new **fc** determined by QPPB. In the ingress datapath SAP ingress QoS policies are applied in the first P chip and route lookup/QPPB occurs in the second P chip. This has the implications listed below:

- Ingress remarking (based on profile state) is always based on the original **fc** (**fc1**) and sub-class (if defined)
- The profile state of a SAP ingress packet that matches a QPPB route depends on the configuration of **fc2** only. If the de-1-out-profile flag is enabled in **fc2** and **fc2** is not mapped to a priority mode queue, then the packet will be marked out of profile if its DE bit = 1. If the profile state of **fc2** is explicitly configured (in or out) and **fc2** is not mapped to a priority mode queue then the packet is assigned this profile state. In both cases, there is no consideration of whether or not **fc1** was mapped to a priority mode queue.
- The priority of a SAP ingress packet that matches a QPPB route depends on several factors. If the de-1-out-profile flag is enabled in **fc2** and the DE bit is set in the packet then priority will be low regardless of the QPPB priority or **fc2** mapping to profile mode queue, priority mode queue or policer. If **fc2** is associated with a profile mode queue then the packet priority will be based on the explicitly configured profile state of **fc2** (in profile = high, out profile = low, undefined = high), regardless of the QPPB priority or **fc1** configuration. If **fc2** is associated with a priority mode queue or policer then the packet priority will be based on QPPB (unless DE=1), but if no priority information is associated with the route then the packet priority will be based on the configuration of **fc1** (if **fc1** mapped to a priority mode queue then it is based on DSCP/IP prec/802.1p and if **fc1** mapped to a profile mode queue then it is based on the profile state of **fc1**).

Table 3 summarizes the interactions.

Table 3 QPPB Interactions with SAP Ingress QoS

Original FC object mapping	New FC object mapping	Profile	Priority (drop preference)	DE=1 override	In/out of profile marking
Profile mode queue	Profile mode queue	From new base FC unless overridden by DE=1	From QPPB, unless packet is marked in or out of profile in which case follows profile. Default is high priority.	From new base FC	From original FC and sub-class

Table 3 QPPB Interactions with SAP Ingress QoS (Continued)

Original FC object mapping	New FC object mapping	Profile	Priority (drop preference)	DE=1 override	In/out of profile marking
Priority mode queue	Priority mode queue	Ignored	If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then from original dot1p/exp/DSCP mapping or policy default.	From new base FC	From original FC and sub-class
Policer	Policer	From new base FC unless overridden by DE=1	If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then from original dot1p/exp/DSCP mapping or policy default.	From new base FC	From original FC and sub-class
Priority mode queue	Policer	From new base FC unless overridden by DE=1	If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then from original dot1p/exp/DSCP mapping or policy default.	From new base FC	From original FC and sub-class
Policer	Priority mode queue	Ignored	If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then from original dot1p/exp/DSCP mapping or policy default.	From new base FC	From original FC and sub-class
Profile mode queue	Priority mode queue	Ignored	If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then follows original FC's profile mode rules.	From new base FC	From original FC and sub-class
Priority mode queue	Profile mode queue	From new base FC unless overridden by DE=1	From QPPB, unless packet is marked in or out of profile in which case follows profile. Default is high priority.	From new base FC	From original FC and sub-class

Table 3 QPPB Interactions with SAP Ingress QoS (Continued)

Original FC object mapping	New FC object mapping	Profile	Priority (drop preference)	DE=1 override	In/out of profile marking
Profile mode queue	Policer	From new base FC unless overridden by DE=1	If DE=1 override, then low otherwise from QPPB. If no DEI or QPPB overrides, then follows original FC's profile mode rules.	From new base FC	From original FC and sub-class
Policer	Profile mode queue	From new base FC unless overridden by DE=1	From QPPB, unless packet is marked in or out of profile in which case follows profile. Default is high priority.	From new base FC	From original FC and sub-class

2.2.1.4 Object Grouping and State Monitoring

This feature introduces a generic operational group object which associates different service endpoints (pseudowires and SAPs) located in the same or in different service instances. The operational group status is derived from the status of the individual components using certain rules specific to the application using the concept. A number of other service entities, the monitoring objects, can be configured to monitor the operational group status and to perform certain actions as a result of status transitions. For example, if the operational group goes down, the monitoring objects will be brought down.

2.2.1.4.1 IES IP Interface Applicability

This concept is used by an IPv4 IES interface to affect the operational state of the IP interface monitoring the operational group. Individual SAP and spoke SDPs are supported as monitoring objects.

The following rules apply:

- an object can only belong to one group at a time
- an object that is part of a group cannot monitor the status of a group
- an object that monitors the status of a group cannot be part of a group
- an operational group may contain any combination of member types: SAP or Spoke-SDPs

- an operational group may contain members from different VPLS service instances
- objects from different services may monitor the oper-group

There are two steps involved in enabling the functionality:

1. Identify a set of objects whose forwarding state should be considered as a whole group then group them under an operational group using the **oper-group** command.
2. Associate the IP interface to the oper-group using the **monitor-group** command.

The status of the operational group (oper-group) is dictated by the status of one or more members according to the following rules:

- The oper-group goes down if all the objects in the oper-group go down. The oper-group comes up if at least one component is up.
- An object in the group is considered down if it is not forwarding traffic in at least one direction. That could be because the operational state is down or the direction is blocked through some validation mechanism.
- If a group is configured but no members are specified yet, then its status is considered up.
- As soon as the first object is configured the status of the operational group is dictated by the status of the provisioned member(s).

The following configuration shows the oper-group g1, the VPLS SAP that is mapped to it and the IP interfaces in IES service 2001 monitoring the oper-group g1. This example uses an R-VPLS context. The VPLS instance includes the **allow-ip-int-bind** and the **name v1**. The IES interface links to the VPLS using the **vpls v1** option. All commands are under the configuration service hierarchy.

To further explain the configuration. Oper-group g1 has a single SAP (1/1/1:2001) mapped to it and the IP interfaces in the IES service 2001 will derive its state from the state of oper-group g1.

```
oper-group g1 create

vpls 1 name "v1" customer 1 create
    allow-ip-int-bind
    stp
        shutdown
    exit
    sap 1/1/1:2001 create
        oper-group g1
        eth-cfm
            mep domain 1 association 1 direction down
ccm-enable
    no shutdown
```



```
        exit
    exit
    sap 1/1/2:2001 create
    exit
    sap 1/1/3:2001 create
    exit
no shutdown

ies 2001 customer 1 create
    interface "i2001" create
        address 192.168.1.1/24
        monitor-oper-group "g1"
        vpls "v1"
    exit
no shutdown
exit
```

2.2.2 Subscriber Interfaces

Subscriber interfaces are composed of a combination of two key technologies, subscriber interfaces and group interfaces. While the subscriber interface defines the subscriber subnets, the group interfaces are responsible for aggregating the SAPs. Subscriber Interfaces apply to the 7450 ESS and 7750 SR only.

- Subscriber interface — An interface that allows the sharing of a subnet among one or many group interfaces in the routed CO model.
- Group interface — Aggregates multiple SAPs on the same port.
- Redundant interfaces — A special spoke-terminated Layer 3 interface. It is used in a Layer 3 routed CO dual-homing configuration to shunt downstream (network to subscriber) to the active node for a given subscriber. Redundant interfaces apply to the 7750 SR only.

2.2.2.1 IPv6 Enhanced Subscriber Management (ESM)

All IPv6 ESM services require either Routed CO (IES), or Routed CO for VPRN as a supporting service construct. Because of the complexities of the IPv6 link-model, there is currently no support for IPv6 ESM in a VPLS. There is also currently no support for IPv6 in combination with Basic Subscriber Management (BSM). This feature applies to the 7450 ESS and 7750 SR only.

2.2.2.2 RADIUS Accounting

In the 7750 SR OS, the accounting paradigm is based on sla-profile instances, yet this is at odds with traditional RADIUS authentication and accounting which is host-centric. In previous OS releases, it was possible to have many hosts sharing a common sla-profile instance, and thus accounting and QoS parameters.

Complications would arise with RADIUS accounting because Accounting-Start and Accounting-Stop are a function of sla-profile instance and not the hosts – this meant that some host-specific parameters (like Framed-Ip-Address) would not be consistently included in RADIUS accounting.

Dual-stack subscribers are now two different hosts sharing a single sla-profile instance. A new RADIUS accounting mode has been introduced to support multiple-host environments.

A new command, **host-accounting**, is introduced under **accounting-policy**, which allows configurable behavior.

No host-accounting:

When **no host-accounting** is configured, accounting behavior is as follows:

- A RADIUS accounting start message is sent when the SLA-profile instance is created. It contains accounting (octets/packets) and the Framed-Ip-Address of the host which caused the sla-profile instance to be created.
- Additional hosts may bind to the sla-profile instance at any time, but no additional accounting messages are sent during these events.
- If the original host disconnects, then future accounting messages will use an IP address of one of the remaining hosts.
- When the final host associated with an sla-profile instance disconnects, an accounting stop message will be sent.

Host-accounting enabled:

When **host-accounting** is configured, additional RADIUS accounting messages are created for host activity in addition to messages for common queue accounting. The behavior is as follows:

- A RADIUS accounting start message is sent each time a host is authenticated. It contains the Framed-Ip-Address among other things. It does not contain any octet or packet counts.
- A RADIUS accounting start message is sent each time a sla-profile instance is created.
- Whenever a host disconnects a RADIUS, accounting stop message is sent for that host.

- If all host associated with an sla-profile instance disconnect, a RADIUS accounting stop message is sent for that instance.

This behavior means certain AVP may be in either host, sla-profile instance, or both accounting records. See [Table 4](#).



Note: Interim-Acct records are not sent for hosts, only the start- and stop-accounting messages.

Table 4 RADIUS Accounting Table

RADIUS Accounting AVP	Host Accounting	SLA-Profile Accounting
User-Name	Yes	No
NAS-Identifier	Yes	Yes
NAS-IP-Address	Yes	Yes
Nas-Port-Id	Yes	No
Nas-Port	Yes	No
Nas-Port-Type	Yes	No
Service-Type	Yes	No
Framed-Protocol	Yes	No
Framed-Ip-Address	Yes	No
Framed-Ip-Netmask	Yes	No
Framed-Route	Yes	No
Class	Yes	No
Session-Timeout	Yes	Yes
Circuit-Id VSA	Yes	No
Called-Station-Id	Yes	No
Calling-Station-Id	Yes	No
MAC-Addr VSA	Yes	No
Remote-Id VSA	Yes	No
Acct-Input-Octets	No	Yes
Acct-Output-Octets	No	Yes

Table 4 RADIUS Accounting Table (Continued)

RADIUS Accounting AVP	Host Accounting	SLA-Profile Accounting
Acct-Input-Gigawords	No	Yes
Acct-Output-Gigawords	No	Yes
Acct-Session-Id	Yes	Yes
Acct-Session-Time	Yes	Yes
Acct-Input-Packets	No	Yes
Acct-Output-Packets	No	Yes
Agent-Circuit-Id	Yes	No
Agent-Remote-Id	Yes	No
Actual-Data-Rate-Upstream	Yes	No
Actual-Data-Rate-Downstream	Yes	No
Access-Loop-Encapsulation	Yes	No
Alc-Accounting	No	Yes
Alc-Subscriber-Id	Yes	Yes
Alc-Subscriber-Profile-String	Yes	Yes
Alc-Sla-Profile-String	Yes	Yes

2.2.3 SAPs

2.2.3.1 Encapsulations

The following SAP encapsulations are supported on the IES services:

- Ethernet null
- Ethernet dot1q
- SONET/SDH IPCP
- SONET/SDH BCP-null
- SONET/SDH BCP-dot1q
- SONET/SDH ATM

- ATM - LLC SNAP or VC-MUX

2.2.3.2 ATM SAP Encapsulations for IES

The 7750 SR series supports ATM PVC service encapsulation for IES SAPs. Both UNI and NNI cell formats are supported. The format is configurable on a SONET/SDH path basis. A path maps to an ATM VC. All VCs on a path must use the same cell format.

The following ATM encapsulation and transport modes are supported:

- RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*:
 - AAL5 LLC/SNAP IPv4 routed
 - AAL5 VC mux IPv4 routed
 - AAL5 LLC/SNAP IPv4 bridged
 - AAL5 VC mux IPv4 bridged

2.2.3.3 Pseudowire SAPs


This feature allows customers of an IES, VPRN, or Epipe VLL service and connected to an Ethernet SAP on an Access PE to be backhauled through an Ethernet aggregation network using MPLS pseudowires terminating directly on a converged PE hosting the IES, VPRN, or Epipe VLL service. If Enhanced Subscriber Management over PW is also used, then the converged PE may also act as a BNG. This service is different from VLL Spoke-SDP termination on an IES or VPRN because access QoS policies can be applied directly at a centralized PE hosting the IES or VPRN instance. This feature uses the same concepts of pseudowire ports and pseudowire SAPs that are used for ESM over MPLS pseudowires, described in the 7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide.

The MPLS pseudowire originates from the first hop aggregation PE (referred to as access PE) upstream of the Access-Node (or directly from a multi-service AN), and terminates on the converged PE. Multiple customers from a given access-port on the Access-PE can be backhauled over a single MPLS pseudowire towards the converged PE. This capability allows the network to scale and does not require an MPLS pseudowire per customer between the Access-PE and the converged PE. The access-port on the Access-PE can be dot1q, q-in-q or NULL encapsulated. The converged PE terminates the MPLS pseudowire, decapsulates the received frames, and provides access QoS functions including HQoS, without requiring an internal or

external loopback. Each MPLS pseudowire is represented on the BNG as a “PW-port” for which SAPs are created. These SAPs are termed “PW SAPs”, and must be statically configured on IES or VPRN interfaces (unlike the ESM case where a capture SAP can be configured). The underlying Ethernet port must be in hybrid mode. Pseudowire SAPs are supported on Ethernet MDAs and on the HSMDAv2.

Figure 3 illustrates the architecture of an aggregation network that uses pseudowire SAPs.

Figure 3 Network Architecture using Pseudowire SAPs



al_0620

2.2.3.4 Encapsulation

The packet is encapsulated on an Ethernet pseudowire, which is associated with a pseudowire port on the converged PE, and a spoke SDP on the access PE. The optional control word is not supported. The SDP could use an LDP LSP, RSVP LSP, segment routed tunnel, BGP RFC3107 tunnel, or LDP over RSVP tunnel. Hash labels are not supported. The SDP may be bound to a port or a LAG, although shaping Vports for pseudowire ports on LAGs in distributed mode is not supported. If an SDP is rerouted, then the corresponding pseudowire ports are brought operationally down. Pseudowire ports are associated with an SDP by configuration.

2.2.3.5 Pseudowire SAP Configuration

The following steps are required at the access PE:

1. Configure an Epipe VLL service.
2. Configure a NULL, 1q or q-in-q SAP on the Epipe service.

The following steps are used to configure a pseudowire SAP on the IES or VPRN service at the Layer 3 PE:

Step 1. Define a pseudowire port.

```
pw-port 1 create
exit
pw-port 2 create
exit
```

Step 2. Bind a physical port or LAG, in hybrid mode, with the pseudowire port.

```
service customer 1 create
    multi-service-site "abc" create
        assignment port pw-1
        egress
        policer-control-policy "abc"
        exit
    exit description "Default customer"
    exit
sdp 1 mpls create
    far-end 10.1.1.2
    ldp
    path-mtu 1514
    keep-alive
    shutdown
    exit
binding
    port lag-1
    pw-port 1 vc-id 1 create
        no shutdown
        exit
    pw-port 2 vc-id 2 create
```

```

        no shutdown
        exit
    exit
no shutdown
exit

```

Step 3. Perform one of the following steps:

Step 4. For a PW SAP on an IES/VPRN, configure a SAP on the IES or VPRN interface, with a SAP ID that uses the form **pw-id**.

```

ies 1 customer 1 create
interface "ies if" create
address 192.168.1.1/24
mac 00:00:00:00:00:ff
static-arp 192.168.1.2 00:00:00:00:00:aa
sap pw-1:1 create
    exit
exit
no shutdown
exit

```

Step 5. For a PW SAP on an Epipe VLL, configure a SAP on the service, with a SAP ID that uses the form **pw-id**.

```

epipe 1 customer 1 create
sap pw-1:1 create
    exit
exit
no shutdown
exit

```

The PW SAP may be mated to an Ethernet SAP or an Ethernet Spoke-SDP in the Epipe VLL service

2.2.3.6 QoS for Pseudowire Ports and Pseudowire SAPs

Pseudowire SAPs support the QoS models allowed for regular VLL, IES or VPRN SAPs. These include:

- Per-service HQoS.

- This allows shaping of the total traffic per access node (and total traffic per class per AN), assuming one pseudowire per AN from the A-PE.
- SAP QoS support as available on the IOM3-XP, including
 - H-QoS (service scheduler child to port scheduler parent)
 - SAP queues attached to H-QoS scheduler by 'parent' statement
 - Scheduler attached to Port Scheduler by 'port-parent' statement
 - Direct service queue to port scheduler mapping
 - Aggregate-rate-limit
 - Support for the redirection of SAP egress queues to an access queue group instance. It is possible to redirect SAP queues of a pseudowire SAP using the SAP based redirection for the IOM with Ethernet MDA or HSMDAv2, and policy based redirection for the IOM with Ethernet MDA, as applicable.
- Policing and H-POL

2.2.3.7 Shaping and Bandwidth Control

Pseudowire SAPs can be shaped on egress by a Vport on a physical port. The pseudowire SAP egress cannot explicitly declare which Vport to use, but they inherit the Vport used by the PW port egress shaping.

The Vport is represented by a secondary shaper on an HSMDAv2. The intermediate destination identifier, used for ESM on MPLS pseudowires, is not applicable to VLL, IES and VPRN pseudowire SAPs.

If a pseudowire port is configured on a LAG, then Vport shaping is only supported if the LAG is in link mode.

Per-access node shaping is configured as follows:

1. Configure a Vport per AN under the port (or LAG) to which the SDP corresponding to the pseudowire SAP is bound. The Vport would be configured with **aggregate rate-limit**. (`config>port>ethernet>access>egress>vport vport-name create`).
2. Explicitly assign (by static configuration) a pseudowire port to a Vport. For limiting the total traffic to an AN, all pseudowire ports for an AN-port would refer to the same Vport.

As in the ESM on pseudowire case, Vport scheduling on the HSMDAv2 is implemented using an **exp-secondary-shaper**. This is referred to as a **pw-sap-secondary-shaper** in the new CLI below. If an **hsmdda-queue-override secondary-shape** is defined for the pw-sap, then the system will use the override, else:


- If a named **pw-sap-secondary-shaper** is defined for the pw-port, then that is used,
- Else, the default **exp-secondary-shaper** for the port is used.

For bandwidth control per pseudowire, the following configuration steps are used:

1. Create multiple Vports under the port to which SDP is bound. Each Vport can be configured with **agg-rate rate**, a scheduler or port-scheduler.
2. Assign each pseudowire to an AN to a unique Vport shaper (regular IOM/MDA) or secondary shaper (on HSMDAv2).

To make use of the **agg-rate rate** or **port-scheduler** under a Vport, PW SAP queues and schedulers must be configured with the **port-parent** command. To make use of a scheduler under a Vport, PW SAP schedulers must be configured with a **parent** command and the **parent-location vport** under the tier 1 of the scheduler policy. The egress hierarchical parenting relationship options are shown in [Figure 4](#). See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Quality of Service Guide Quality of Service guide* for more information.

Figure 4 PW SAP Egress Scheduling Hierarchy Options



al_0376

2.2.3.8 Lag Considerations

PW ports may be bound to Vport schedulers bound to a LAG. However, if the LAG is configured in distributed mode, then bandwidth is shared according to the active LAG members across a single IOM. If the LAG spans multiple IOMs, then it effectively operates in link mode across the IOMs. That is, the full LAG bandwidth is allocated to the LAG members on each IOM. Therefore, the use of a Vport on a distributed mode LAG with a port scheduler on the port or Vport and PW SAPs is explicitly not supported and is not a recommended configuration. It is recommended that port-fair mode is used instead.

2.2.3.9 Last Mile Packet Size Adjustment

In the application where pseudowire SAPs are used to apply access QoS for services aggregated from an Ethernet access network, MPLS labels may not be present on the last-mile and link from an access node. In these cases, policers, queues and H-QoS schedulers should account for packets without MPLS overhead, modeled as “encaps-offset”. Vport and port schedulers behave as per the table below. In the data-path, the actual pseudowire encap overhead (taking into account the MPLS labels) added to the packet is tracked, and may be applied to the scheduler calculations via the configured packet-byte-offset.

The exp-secondary-shaper used on the HSMDAv2 always assumes MPLS overhead and does not account for the packet-byte-offset. In all other cases, the rate limit configured for the pseudowire SAP accounts for subscriber or service frame wire rate: without MPLS overhead and including the last mile overhead (unless a packet-byte-offset is configured).

[Table 5](#) summarizes the default packet sizes used at each of the schedulers on the IOM/Ethernet MDA and HSMDAv2, assuming a 1000byte customer packet.

Table 5 Packet Sizes Used for Pseudowire SAPs

Type	Size
exp-secondary-shaper	20B preamble + 26 MPLS + 1000B pkt
queue/policer rate on hsmдав2	1000B customer pkt
port-scheduler rate	20B preamble + 1000B pkt
regular queue/policer rate	1000B pkt
vport agg-limit-rate	20B preamble + 1000B pkt
vport port-scheduler rate	20B preamble + 1000B pkt


Table 5 Packet Sizes Used for Pseudowire SAPs (Continued)

Type	Size
vport scheduler rate	1000B pkt
vport scheduler to port-scheduler rates	20B preamble + 1000B pkt

2.2.3.10 Redundancy with Pseudowire SAPs

Within a chassis, IOM and port based redundancy is based on active/backup LAG. The topology for the base MPLS LSP used by the SDP could be constrained such that it could get re-routed in the aggregation network, but would always appear on the LAG ports on the Layer 3 PE. In the case that the tunnel is re-routed to a different port, the MPLS pseudowire SAPs would be brought down.

In order to provide Layer 3 PE redundancy, dual homing of the access PE into separate Layer 3 PEs using active/standby pseudowire status is supported. This is shown in [Figure 5](#).

Figure 5 Dual Homing into Multiple Layer 3 PEs

Dual homing operates in a similar manner to Spoke-SDP termination on IES/VPRN. [Figure 5](#) displays the access PE is dual-homed to the Layer 3 PEs using two spoke-SDPs. The endpoint in the access PE is configured to be the master from a pseudowire redundancy perspective using the **standby-signaling-master** command. The access PE picks one of the Spoke-SDPs to make active, and one to make standby, based on the local configuration of primary or Spoke SDP precedence.


The pseudowire port at the Layer 3 PE behaves as a slave from the perspective of pseudowire status signaling. That is, if its peer signals “PW FWD standby (0x20)” status bit for the given Spoke-SDP and the local configuration does not allow this bit to be ignored, the PE will take the pseudowire port to a local operationally down state. This is consistent with the Spoke-SDP behavior for the case of Spoke-SDP termination on IES/VRN.

As a consequence, all of the pseudowire SAPs bound to the pseudowire port are taken down, which causes the corresponding IES or VPRN interface to go to a local operationally down state and thus will stop forwarding packets towards this pseudowire port.

Conversely, the formerly standby pseudowire is made active and then the corresponding pseudowire port on the second Layer 3 PE is taken locally operationally up. Therefore, all of the pseudowire SAPs bound to the pseudowire port are brought up, which causes the corresponding IES or VPRN interface to go to a local operationally up state allowing forwarding of packets towards this pseudowire port.

For VLLs, a PW port always behaves as a slave from the perspective of PW redundancy. This is because the PW port is taken locally operationally down if any non-zero PW status (including a PW Preferential Forwarding status of **standby**) is received. Support for existing master-slave PW redundancy mechanisms for dual homing of the access PE into separate converged PEs using active/standby PW status is required as shown in [Figure 6](#).

Figure 6 Master-Slave PW Redundancy



As in the existing implementation, standby-signaling-master is configured on the Spoke-SDP at the access PE. However explicit configuration of standby-signaling-slave on the PW port is not required, as this is the default behavior.

The forwarding behavior is the same as when standby-signaling-slave for Epipe Spoke-SDPs. That is, when enabled, if a PW Forwarding Standby (0x20) LDP status message is received for the PW, then the transmit direction is blocked for the PW port. All PW SAPs bound to the corresponding PW port are treated from a SAP OAM perspective in the same manner as a fault on the service, such as an SDP-binding down or remote SAP down.

PW redundancy with multiple active/standby PW ports or PW SAPs bound to the same Ethernet SAP in the converged PE is not supported. The independent mode of operation for PW redundancy is not also supported for a PW port.

2.2.3.11 Operational Group Support for PW Ports

A PW port state may be linked to the state of an oper-group, such that if the oper-group goes down, the SDP binding for the PW port will also go operationally down, and thus the corresponding PW status bit signaled (0x00000001 - Pseudowire Not Forwarding). If a status of 0x00000001 is signaled for a currently active PW, and active/standby dual homing is in use then the access PE will fail over to the standby PW to the standby converged PE.

This is achieved by linking an SDP binding to an operational group for PW SAPs belonging to any supported service types (including those with group interfaces) bound to that PW port, such as IES, VPRN, or Epipe VLL. The association to an operational group is configured under the PW port config at the SDP binding level, as follows:

```
config
  service
    sdp
      binding
        [no] pw-port <pw-port-id> [vc-id <vc-id>] [create]
        monitor-oper-group <group-name>
```

The **monitor-oper-group** command specifies the operational group to be monitored by the PW-Port under which it is configured. The oper-group name must be already configured under the **config>service** context before its name is referenced in this command.

The following illustrates how a PW port can track the status of VPRN uplinks using monitor-oper-group.

Uplinks in a VPRN may be monitored using a BFD session on the network facing IP interfaces in a VPRN or on the network IP interfaces supporting the uplinks.

Oper-groups monitor the state of these BFD sessions inside the VPRN as follows:

```
config>service>  
  oper-group "test-oper-grp" create  
    bfd-enable interface "vprn-if" dest-ip 10.0.0.20 service 105
```

Alternatively, the state of network interfaces can be monitored as follows:

```
config>service>  
  oper-group "test-oper-grp" create  
    bfd-enable interface "network-if" dest-ip 10.0.1.20
```

The PW port is then configured with monitor-oper-group as follows:

```
config>service>sdp>binding  
  pw-port 100 vc-id 25  
  monitor-oper-group "test-oper-group"
```

2.2.4 Routing Protocols

The IES IP interfaces are restricted as to the routing protocols that can be defined on the interface based on the fact that the customer has a different routing domain for this service. The IES IP interfaces support the following routing protocols:

- RIP
- OSPF
- IS-IS
- BGP
- IGMP
- PIM



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.

2.2.4.1 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.

2.2.5 QoS Policies

When applied to IES services, service ingress QoS policies only create the unicast queues defined in the policy. The multipoint 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. Both Layer 2 or Layer 3 criteria can be used in the QoS policies for traffic classification in an IES.

2.2.6 Filter Policies

Only IP filter policies can be applied to IES services.

2.2.7 MPLS Entropy Label and Hash Label

The router supports both the MPLS entropy label (RFC 6790) and the Flow Aware Transport label, known as the hash label (RFC 6391). These labels allow LSR nodes in a network to load-balance labeled packets in a much more granular fashion than allowed by simply hashing on the standard label stack. See the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR MPLS Guide* for further information.

2.2.8 Spoke SDPs


Distributed services use service distribution points (SDPs) to direct traffic to another router through service tunnels. SDPs are created on each participating router and then bound to a specific service. SDP can be created as either GRE or MPLS. Refer to the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Services Overview Guide* for information about configuring SDPs.

This feature provides the ability to cross-connect traffic entering on a spoke SDP, used for Layer 2 services (VLLs or VPLS), on to an IES or VPRN service. From a logical point of view, the spoke SDP entering on a network port is cross-connected to the Layer 3 service as if it entered by a service SAP. The main exception to this is traffic entering the Layer 3 service by a spoke SDP is handled with network QoS policies not access QoS policies.

This feature applies to the 7450 ESS and 7750 SR only.


Figure 7 depicts traffic terminating on a specific IES or VPRN service that is identified by the SDP-ID and VC label present in the service packet.

Figure 7 SDP-ID and VC Label Service Identifiers



al_0163

Figure 8 IES Spoke-SDP Termination



OSSG188

Figure 8 depicts a spoke-SDP terminating directly into a Layer 3 service interface (IES or VPRN) at one end, and a Layer 2 service (Epipe, lpipe, or VPLS) at the other. There is no special configuration required on the Layer 2 service.

If the terminating Layer 2 service is an Ipipe, then on the IES/VPRN interface end, the spoke-SDP must be created with the vc-type ipipe option. Spoke-SDPs created with vc-type ether (the default) are compatible with Epipe and VPLS services, as well as with other IES/VPRN interfaces.

If the MPLS network uses LDP signaling, then in order for a spoke-SDP to function, the LDP binding MTUs at each end must match. For a Layer 2 service, the MTU of the local binding is 14 octets less than the configured service-mtu (such as, binding MTU = service-mtu - 14). For an IES or VPRN interface, the binding MTU is equal to either the configured ip-mtu of the interface, or the SDP's path-mtu minus 14, whichever is lower. The local and remote MTUs of all bindings can be found using the CLI command **show router ldp bindings**.

All routing protocols that are supported by IES/VPRN are supported for spoke-SDP termination.

Refer to “VCCV BFD support for VLL, Spoke SDP Termination on IES and VPRN, and VPLS Services” in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for information about using VCCV BFD in spoke-SDP termination.



Note: Spoke-SDP termination of Ipipe VLLs on IES is not supported in System Profile B. To determine if Ipipes are currently bound to an IES interface, use the **show router ldp bindings services** command before configuring profile B.

2.2.9 SRRP

Subscriber Router Redundancy Protocol (SRRP) is used on the 7750 SR and 7450 ESS and is closely tied to the multi-chassis synchronization (MCS) protocol used to synchronize information between redundant nodes. An MCS peer must be configured and operational when subscriber hosts have a redundant connection to two nodes. Subscriber hosts are identified by the ingress SAP, the host's IP and MAC addresses. Once a host is identified on one node, the MCS peering is used to inform the other node that the host exists and conveys the dynamic DHCP lease state information of the host. MCS creates a common association between the virtual ports (SAPs) shared by a subscriber. This association is configured at the MCS peering level by defining a tag for a port and range of SAPs. The same tag is defined on the other nodes peering context for another port (does not need to be the same port-ID) with the same SAP range. In this manner, a subscriber host and Dot1Q tag sent across the peering with the appropriate tag is mapped to the redundant SAP on the other node.

SRRP can only be configured on group interfaces. Once SRRP is active on a group IP interface, the SRRP instance attempts to communicate through in-band (over the group IP interfaces SAPs) and out-of-band (over the group IP interfaces redundant IP interface) messages to a remote router. If the remote router is also running SRRP with the same SRRP instance ID, one router enters a master state while the other router enters a backup state. Since both routers are sharing a common SRRP gateway MAC address that is used for the SRRP gateway IP addresses and for proxy ARP functions, either node may act as the default gateway for the attached subscriber hosts.

For proper operation, each subscriber subnet associated with the SRRP instance must have a gw-address defined. The SRRP instance cannot be activated (no shutdown) unless each subscriber subnet associated with the group IP interface has an SRRP gateway IP address. Once the SRRP instance is activated, new subscriber subnets cannot be added without a corresponding SRRP gateway IP address. [Table 6](#) describes how the SRRP instance state is used to manage access to subscriber hosts associated with the group IP interface.

SRRP instances are created in the disabled state (shutdown). To activate SRRP the **no shutdown** command in the SRRP context must be executed.

Before activating an SRRP instance on a group IP interface, the following actions are required:

- Add a SRRP gateway IP addresses to all subscriber subnets associated with the group IP interface, including subnets on subscriber IP interfaces associated as retail routing contexts (at least one subnet must be on the subscriber IP interface containing the group IP interface and its SRRP instance).
- Create a redundant IP interface and associate it with the SRRP instances group IP interface for shunting traffic to the remote router when master.
- Specify the group IP interface SAP used for SRRP advertisement and Information messaging.

Before activating an SRRP instance on a group IP interface, the following actions should be considered:

- Associate the SRRP instance to a Multi-Chassis Synchronization (MCS) peering terminating on the neighboring router (the MCS peering should exist as the peering is required for redundant subscriber host management).
- Define a description string for the SRRP instance.
- Specify the SRRP gateway MAC address used by the SRRP instance (must be the same on both the local and remote SRRP instance participating in the same SRRP context).
- Change the base priority for the SRRP instance.

- Specify one or more VRRP policies to dynamically manage the SRRP instance base priority.
- Specify a new keep alive interval for the SRRP instance.

[Table 6](#) lists the SRRP's state effect on subscriber hosts associated with group IP interfaces.

Table 6 SRRP State Effect on Subscriber Hosts Associated with Group IP Interface

SRRP State	ARP	Local Proxy ARP Enabled	Remote Proxy ARP Enabled	Subscriber Host Routing
Disabled	- Responds to ARP for all owned subscriber subnet IP addresses. - Will not respond to ARP for subscriber subnet SRRP gateway IP addresses. - All ARP responses will contain the native MAC of the group IP interface (not the SRRP gateway MAC).	- Responds to ARP for all subscriber hosts on the subscriber subnet.	- Responds to ARP for all reachable remote IP hosts.	- All routing out the group IP interface will use the native group IP interface MAC address. - The group IP interface redundant IP interface will not be used. - Will not accept packets destined to the SRRP gateway MAC received on the group IP interface.
Becoming Master (In order to enter becoming master state, a master must currently exist)	- Responds to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC). - Responds to ARP for subscriber subnet SRRP gateway IP addresses. (hardware address = SRRP gateway IP address, source MAC = group IP interface native MAC).	- Responds to ARP for all subscriber hosts on the subscriber subnet (hardware address = SRRP gateway MAC address, source MAC = group IP interface native MAC).	- Responds to ARP for all reachable remote IP hosts (hardware address = SRRP gateway MAC address, source MAC = group IP interface native MAC).	- All routing out the group IP interface use the native group IP interface MAC address. - Subscriber hosts mapped to the redundant IP interface are remapped to the group IP interface. - Will not accept packets destined to the SRRP gateway MAC received on the group IP interface.

Table 6 SRRP State Effect on Subscriber Hosts Associated with Group IP Interface

SRRP State	ARP	Local Proxy ARP Enabled	Remote Proxy ARP Enabled	Subscriber Host Routing
Master	- Responds to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC). - Responds to ARP for subscriber subnet SRRP gateway IP addresses (hardware address = SRRP gateway IP address, source MAC = group IP interface native MAC).	- Responds to ARP for all subscriber hosts on the subscriber subnet (hardware address = SRRP gateway MAC address, source MAC = group IP interface native MAC).	- Responds to ARP for all reachable remote IP hosts (hardware address = SRRP gateway MAC address, source MAC = group IP interface native MAC).	- All routing out the group IP interface will use the SRRP gateway MAC address. - Subscriber hosts mapped to the redundant IP interface are remapped to the group IP interface. - Will accept packets destined to the SRRP gateway MAC received on the group IP interface.
Becoming Backup (redundant IP interface operational)	- Responds to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC). - Will not respond to ARP for subscriber subnet SRRP gateway IP addresses.	- Will not respond to ARP for any subscriber hosts on the subscriber subnet.	- Will not respond to ARP for any remote IP hosts.	- Will not route out the group IP interface for subscriber hosts associated with the subscriber subnet. - Will accept packets destined to the SRRP gateway MAC received on the group IP interface. - Subscriber hosts mapped to the group IP interface are remapped to the redundant IP interface.

Table 6 SRRP State Effect on Subscriber Hosts Associated with Group IP Interface

SRRP State	ARP	Local Proxy ARP Enabled	Remote Proxy ARP Enabled	Subscriber Host Routing
Becoming Backup (redundant IP interface not available)	- Responds to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC). - Will not respond to ARP for subscriber subnet SRRP gateway IP addresses.	- Will not respond to ARP for any subscriber hosts on the subscriber subnet.	- Will not respond to ARP for any remote IP hosts.	- Will route out the group IP interface for subscriber hosts associated with the subscriber subnet using the group IP interface native MAC address. - Subscriber hosts mapped to the redundant IP interface are remapped to the group IP interface. - Will accept packets destined to the SRRP gateway MAC received on the group IP interface.
Backup (redundant IP interface operational)	- Responds to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC). - Will not respond to ARP for subscriber subnet SRRP gateway IP addresses.	- Will not respond to ARP for any subscriber hosts on the subscriber subnet.	- Will not respond to ARP for any remote IP hosts.	- Will not route out the group IP interface for subscriber hosts associated with the subscriber subnet. - Subscriber hosts mapped to the group IP interface are remapped to the redundant IP interface. - Will not accept packets destined to the SRRP gateway MAC received on the group IP interface.

Table 6 SRRP State Effect on Subscriber Hosts Associated with Group IP Interface

SRRP State	ARP	Local Proxy ARP Enabled	Remote Proxy ARP Enabled	Subscriber Host Routing
Backup (redundant IP interface not available)	- Responds to ARP for all owned subscriber subnet IP addresses (hardware address and source MAC = group IP interface native MAC). - Will not respond to ARP for subscriber subnet SRRP gateway IP addresses.	- Will not respond to ARP for any subscriber hosts on the subscriber subnet.	- Will not respond to ARP for any remote IP hosts.	- Will route out the group IP interface for subscriber hosts associated with the subscriber subnet using the group IP interface native MAC address. - Subscriber hosts mapped to the redundant IP interface are remapped to the group IP interface. - Will not accept packets destined to the SRRP gateway MAC received on the group IP interface.

2.2.9.1 SRRP Messaging

SRRP uses the same messaging format as VRRP with slight modifications. The source IP address is derived from the system IP address assigned to the local router. The destination IP address and IP protocol are the same as VRRP (224.0.0.18 and 112, respectively).

The message type field is set to 1 (advertisement) and the protocol version is set to 8 to differentiate SRRP message processing from VRRP message processing.

The vr-id field supports an SRRP instance ID of 32 bits.

Due to the large number of subnets backed up by SRRP, only one message every minute carries the gateway IP addresses associated with the SRRP instance. These gateway addresses are stored by the local SRRP instance and are compared with the gateway addresses associated with the local subscriber IP interface.

Unlike VRRP, only two nodes may participate in an SRRP instance due the explicit association between the SRRP instance group IP interface, the associated redundant IP interface and the multi-chassis synchronization (MCS) peering. Since only two nodes are participating, the VRRP skew timer is not utilized when waiting to enter the master state. Also, SRRP always preempts when the local priority is better than the current master and the backup SRRP instance always inherits the master's advertisement interval from the SRRP advertisement messaging.

SRRP advertisement messages carry a *becoming-master* indicator flag. The *becoming-master* flag is set by a node that is attempting to usurp the master state from an existing SRRP master router. When receiving an SRRP advertisement message with a better priority and with the *becoming-master* flag set, the local master initiates its *becoming-backup* state, stops routing with the SRRP gateway MAC and sends an SRRP advertisement message with a priority set to zero. The new master continues to send SRRP advertisement messages with the *becoming-master* flag set until it either receives a return priority zero SRRP advertisement message from the previous master or its *becoming-master* state timer expires. The new backup node continues to send zero priority SRRP advertisement messages every time it receives an SRRP advertisement message with the *becoming-master* flag set. After the new master either receives the old masters priority zero SRRP advertisement message or the *become-master* state timer expires, it enters the *master* state. The *become-master* state timer is set to 10 seconds upon entering the *become-master* state.

The SRRP advertisement message is always evaluated to see if it has a higher priority than the SRRP advertisement that would be sent by the local node. If the advertised priority is equal to the current local priority, the source IP address of the received SRRP advertisement is used as a tie breaker. The node with the lowest IP address is considered to have the highest priority. SRRP will not preempt when priorities are equal. Preemption occurs only when priorities are specified. The lower IP address is only used as a tie-breaker when there is no master in the network. In other words, when both routers are changing from the "init" state, the lower IP will be used to choose the master. If a master already exists, despite having the lower IP address, the system will not preempt the current master.

The SRRP instance maintains the source IP address of the current master. If an advertisement is received with the current masters source IP address and the local priority is higher priority than the masters advertised priority, the local node immediately enters the *becoming-master* state unless the advertised priority is zero. If the advertised priority is zero, the local node bypasses the *becoming-master* state and immediately enters the *master* state. Priority zero is a special case and is sent when an SRRP instance is relinquishing the master state.

2.2.9.2 SRRP and Multi-Chassis Synchronization

In order to take full advantage of SRRP resiliency and diagnostic capabilities, the SRRP instance should be tied to a MCS peering that terminates on the redundant node. The SRRP instance is tied to the peering using the **srrp srrp-id** command within the appropriate MCS peering configuration. Once the peering is associated with the SRRP instance, MCS will synchronize the local information about the SRRP instance with the neighbor router. MCS automatically derives the MCS key for the SRRP instance based on the SRRP instance ID. For example, an SRRP instance ID of 1 would appear in the MCS peering database with a MCS-key srrp-0000000001.

The SRRP instance information stored and sent to the neighbor router consists of:

- The SRRP instance MCS key
- Containing service type and ID
- Containing subscriber IP interface name
- Subscriber subnet information
- Containing group IP interface information
- The SRRP group IP interface redundant IP interface name, IP address and mask
- The SRRP advertisement message SAP
- The local system IP address (SRRP advertisement message source IP address)
- The Group IP interface MAC address
- The SRRP gateway MAC address
- The SRRP instance administration state (up / down)
- The SRRP instance operational state (disabled / becoming-backup / backup / becoming-master / master)
- The current SRRP priority
- Remote redundant IP interface availability (available / unavailable)
- Local receive SRRP advertisement SAP availability (available / unavailable)

2.2.9.3 SRRP Instance

The SRRP instance uses the received information to verify provisioning and obtain operational status of the SRRP instance on the neighboring router.

2.2.9.3.1 SRRP Instance MCS Key

The SRRP instance MCS key ties the received MCS information to the local SRRP instance with the same MCS key. If the received key does not match an existing SRRP instance, the MCS information associated with the key is ignored. Once an SRRP instance is created and mapped to an MCS peering, the SRRP instance evaluates received information with the same MCS key to verify it corresponds to the same peering. If the received MCS key is on a different peering than the local MCS key an SRRP peering mismatch event is generated detailing the SRRP instance ID, the IP address of the peering the MCS key is received on and the IP address to which the local MCS key is mapped. If the peering association mismatch is corrected, an SRRP peering mismatch clear event is generated.

2.2.9.3.2 Containing Service Type and ID

The Containing Service Type is the service type (IES or VPRN) that contains the local SRRP instance. The Containing Service ID is the service ID of that service. This information is supplied for troubleshooting purposes only and is not required to be the same on both nodes.

2.2.9.3.3 Containing Subscriber IP Interface Name

The containing subscriber IP interface name is the subscriber IP interface name that contains the SRRP instance and its group IP interface. This information is supplied for troubleshooting purposes only and is not required to be the same on both nodes.

2.2.9.3.4 Subscriber Subnet Information

The subscriber subnet information includes all subscriber subnets backed up by the SRRP instance. The information for each subnet includes the Owned IP address, the mask and the gateway IP address. If the received subscriber subnet information does not match the local subscriber subnet information, an SRRP Subscriber Subnet Mismatch event is generated describing the SRRP instance ID and the local and remote node IP addresses. Once the subscriber subnet information matches, an SRRP Subscriber Subnet Mismatch Clear event is generated.

2.2.9.3.5 Containing Group IP Interface Information

The containing group IP interface information is the information about the group IP interface that contains the SRRP instance. The information includes the name of the group IP interface, the list of all SAPs created on the group IP interface, the administrative and operational state of each SAP and the MCS key and the peering destination IP address associated with each SAP. To obtain the MCS information, the SRRP instance queries MCS to determine the peering association of the SRRP instance and then queries MCS for each SAP on the group IP interface. If the local SRRP instance is associated with a different MCS peering than any of the SAPs or if one or more SAPs are not tied to an MCS peering, an SRRP group interface SAP peering mismatch event is generated detailing the SRRP instance ID, and the group IP interface name.

When receiving the remote containing group IP interface information, the local node compares the received SAP information with the local group IP interface SAP information. If a local SAP is not included in the SAP information or a remote SAP is not included in the local group IP interface, an SRRP Remote SAP mismatch event is generated detailing the SRRP instance ID and the local and remote group IP interface names. If a received SAP's MCS key does not match a local SAP's MCS Key, an SRRP SAP MCS key mismatch event is generated detailing the SRRP instance ID, the local and remote group IP interface names, the SAP-ID and the local and remote MCS keys.

2.2.9.3.6 Remote Redundant IP Interface Mismatch

If the group IP remote redundant IP interface address space does not exist, is not within the local routing context for the SRRP instances group IP interface or is not on a redundant IP interface, the local node sends redundant IP interface unavailable to prevent the remote neighbor from using its redundant IP interface. An SRRP redundant IP interface mismatch event is generated for the SRRP instance detailing the SRRP instance, the local and remote system IP addresses, the local and remote group IP interface names and the local and remote redundant IP interface names and IP addresses and masks. The local redundant IP interface may still be used if the remote node is not sending redundant IP interface unavailable.

2.2.9.3.7 Remote Sending Redundant IP Interface Unavailable

If the remote node is sending redundant IP interface unavailable, the local node will treat the local redundant IP interface associated with the SRRP instances group IP interface as down. A Local Redundant IP Interface Unavailable event is generated detailing the SRRP instance ID, the local and remote system IP addresses, the local group IP interface name, the local redundant IP interface name and the redundant IP interface IP address and mask.

2.2.9.3.8 Remote SRRP Advertisement SAP Non-existent

If the remote node's SRRP advertisement SAP does not exist on the local SRRP instances group IP interface, the local node sends local receive SRRP advertisement SAP unavailable to the remote node. An SRRP receive advertisement SAP non-existent event is generated detailing the SRRP instance ID, the local and remote system IP addresses, the local group IP interface name and the received remote SRRP advertisement SAP. Since SRRP advertisement messages cannot be received, the local node will immediately become master if it has the lower system IP address.

2.2.9.3.9 Remote Sending Local Receive SRRP Advertisement SAP Unavailable

If the local node is receiving local receive SRRP advertisement SAP unavailable from the remote node, an SRRP Remote Receive advertisement SAP Unavailable event will be generated detailing the SRRP instance ID, the local and remote system IP addresses, the remote group IP interface name and the local SRRP advertisement SAP. Since the remote node cannot receive SRRP advertisement messages, the local node will immediately become master if it has the lower system IP address.

2.2.9.3.10 Local and Remote Dual Master Detected

If the local SRRP state is master and the remote SRRP state is master, an SRRP dual master event is generated detailing the SRRP instance ID and the local, remote system IP addresses and the local and remote group IP interface names and port numbers.

2.2.9.4 Subscriber Subnet Owned IP Address Connectivity

In order for the network to reliably reach the owned IP addresses on a subscriber subnet, it is not necessary for the owning node to advertise the IP addresses as /32 host routes into the core. Network reachability to the subscriber subnet is advertised into the IGP core by both of the dual homing nodes. The shortest path to the subscriber may not always traverse the active path for a subscriber. In this case, the path traverses the non-active/primary node for the subscriber and the traffic will be redirected through the redundant interface to the other node through the redundant interface to the active path. This ensures that all downstream traffic to a given subscriber will always flow through one node.

2.2.9.5 Subscriber Subnet SRRP Gateway IP Address Connectivity

The SRRP gateway IP addresses on the subscriber subnets cannot be advertised as /32 host routes since they may be active (master) on multiple group IP interfaces on multiple SRRP routers. Without a /32 host route path, the network will forward any packet destined to an SRRP gateway IP address to the closest router advertising the subscriber subnet. While a case may be made that only a node that is currently forwarding for the gateway IP address in a master state should respond to ping or other diagnostic messages, the distribution of the subnet and the case of multiple masters make any resulting response or non-response inconclusive at best. To provide some ability to ping the SRRP gateway address from the network side reliably, any node receiving the ICMP ping request responds if the gateway IP address is defined on its subscriber subnet.

2.2.9.6 Receive SRRP Advertisement SAP and Anti-Spoof

The group IP interface SAPs are designed to support subscriber hosts and perform an ingress anti-spoof function that ensures that any IP packet received on the group IP interface is coming in the correct SAP with the correct MAC address. If the IP and MAC are not registered as valid subscriber hosts on the SAP, the packet is silently discarded. Since the SRRP advertisement source IP addresses are not subscriber hosts, an anti-spoof entry will not exist and SRRP advertisement messages would normally be silently discarded. To avoid this issue, when a group IP interface SAP is configured to send and receive SRRP advertisement messages, anti-spoof processing on the SAP is disabled. This precludes subscriber host management on the SRRP messaging SAP.

2.2.9.7 BFD with SRRP/VRRP

BFD with SRRP is supported. This allows the use of longer timers inside SRRP resulting in more SRRP instances while still retaining fast failure detection with BFD.

2.3 Configuring an IES Service with CLI

This section provides information to configure IES services using the command line interface.

2.3.1 Basic Configuration

The most basic IES service configuration has the following entities:

- customer ID (refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Services Overview Guide* for more information)
- an interface to create and maintain IP routing interfaces within IES service ID
- a SAP on the interface specifying the access port and encapsulation values

The following is a sample configuration of an IES service on ALA-48:

```
*A:ALA-48>config>service# info
-----
ies 1000 customer 50 vpn 1000 create
    description "to internet"
    interface "to-web" create
        address 10.1.1.1/24
        sap 1/1/5:0.* create
    exit
    exit
    no shutdown
-----
*A:ALA-48>config>service#
```

2.3.2 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.

- Step 1.** Associate an IES service with a customer ID.
- Step 2.** Associate customer ID with the service.
- Step 3.** Assign an IP address.
- Step 4.** Create a subscriber interface (for 7450 ESS and 7750 SR only and optional).
- Step 5.** Create an interface.
- Step 6.** Define SAP parameters on the interface

- Select node(s) and port(s).
- Optional — select QoS policies other than the default (configured in the **config>qos** context); for 7450 ESS and 7750 SR only.
- Optional — select filter policies (configured in the **config>filter** context).
- Optional — select accounting policy (configured in the **config>log** context); for 7450 ESS and 7750 SR only.

Step 7. Enable service.

2.3.3 Configuring IES Components

Use the CLI syntax to configure IES components.

2.3.3.1 Configuring an IES Service

Use the following CLI syntax to create an IES service:

The following example displays a basic IES service configuration for the 7450 ESS and 7750 SR:

```
A:ALA-48>config>service#
-----
...
ies 1001 customer 1730 vpn 1001 create
    description "to-internet"
    no shutdown
exit
-----
A:ALA-48>config>service#
```

The following example displays a basic IES service configuration for the 7950 XRS:

```
A:ALA-48>config>service#
-----
...
ies 1001 customer 1730 create
    description "to-internet"
    no shutdown
exit
-----
A:ALA-48>config>service#
```


2.3.3.2 Configuring IES Subscriber Interface Parameters



Note: This section applies only to the 7750 SR only.

Subscriber interfaces operate only with basic (or enhanced) subscriber management. At the very least, a host, either statically configured or dynamically learned by DHCP must be present in order for the interface to be useful.

Refer to [IES Services Command Reference](#) for CLI syntax to configure IES subscriber interface parameters.

The following example displays a subscriber interface configuration for the 7750 SR:

```
A:ALA-48>config>service>ies>sub-if# info
-----
address 192.168.140.1/24
group-interface "abc-if" create
sap 1/1/19:0 create
  ingress
    qos 2
    filter ip 10
  exit
static-host ip 192.168.145.100 mac 00:01:00:00:00:01 create
exit
exit
exit
-----
A:ALA-48>config>service>ies>sub-if#
```

2.3.3.3 Configuring IES Interface Parameters

The following example displays an IES configuration with interface parameters for the 7450 ESS and 7750 SR:

```
A:ALA-48>config>service>ies>if# info
-----
address 10.1.1.1/24
sap 1/1/10:0.* create
  ingress
    qos 100
  exit
  egress
    scheduler-policy "SLA1"
  exit
exit
vrrp 1 owner
authentication-key "3WErEDoZxyQ" hash
```

```
exit
-----
A:ALA-48>config>service>ies>if#
```

The following example displays an IES configuration with interface parameters for the 7950 XRS:

```
A:ALA-48>config>service>ies>if# info
-----
address 10.1.1.1/24
sap 1/1/10:0.* create
exit
-----
A:ALA-48>config>service>ies>if#
```

2.3.3.4 Configuring Spoke-SDP Parameters

The following example displays a spoke SDP configuration for the 7450 ESS and 7750 SR:

```
A:ALA-48>config>service>ies# info
-----
description "to internet"
interface "spokeSDP-test" create
spoke-sdp 2:100 create
egress
filter ip 10
exit
exit
no shutdown
-----
A:ALA-48>config>service>ies#
```

2.3.3.5 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 SAP parameters on a 7450 ESS or 7750 SR, 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.

IES interface ATM SAP parameters can only be configured on ATM-type MDAs and ATM-configured ports of the 7750 SR. See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide*.

Refer to [IES Services Command Reference](#) for CLI syntax.

This example displays an IES SAP configuration for the 7450 ESS and 7750 SR:

```
*A:ALA-A>config>service>ies>if# info
-----
address 10.10.36.2/24
sap 5/1/3.1:0 create
  ingress
    qos 101
  exit
  egress
    scheduler-policy "alpha"
    qos 1010
  exit
exit
-----
*A:ALA-A>config>service>ies>if#
```

This example displays an IES SAP configuration for the 7950 XRS:

```
*A:ALA-A>config>service>ies>if# info
-----
address 10.10.36.2/24
exit
-----
*A:ALA-A>config>service>ies>if#
```

2.3.3.6 Configuring IES SAP ATM Parameters



Note: This section applies only to the 7750 SR only.

The following example displays the command usage to create Apipe SAPs:

PE router 1 (A:ALA-41):

```
Example: A:ALA-41>config>service# ies 5
A:ALA-41>config>service>ies# sap 1/1/1:0/32 create
A:ALA-41>config>service>ies>sap# ingress
A:ALA-41>config>service>ies>sap>ingress# qos 102
A:ALA-41>config>service>ies>sap>ingress# exit
A:ALA-41>config>service>ies>sap# egress
```

```
A:ALA-41>config>service>ies>sap>egress# qos 103
A:ALA-41>config>service>ies>sap>egress# exit
A:ALA-41>config>service>ies>sap# no shutdown
A:ALA-41>config>service>ies>sap# exit
A:ALA-41>config>service>ies#
```

PE router 2 (A:ALA-42):

Example:

```
A:ALA-42>config>service# ies 5
A:ALA-42>config>service>ies# sap 2/2/2:0/32 create
A:ALA-42>config>service>ies>sap# ingress
A:ALA-42>config>service>ies>sap>ingress# qos 102
A:ALA-42>config>service>ies>sap>ingress# exit
A:ALA-42>config>service>ies>sap# egress
A:ALA-42>config>service>ies>sap>egress# qos 103
A:ALA-42>config>service>ies>sap>egress# exit
A:ALA-42>config>service>ies>sap# no shutdown
A:ALA-42>config>service>ies>sap# exit
A:ALA-42>config>service>ies#
```

The following output displays the IES SAP configuration.

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
-----
...
    ies 5 customer 1 create
        description "apipe test"
        service-mtu 1400
        sap 1/1/1:0/32 create
            ingress
                qos 102
            exit
            egress
                qos 103
            exit
        exit
        no shutdown
    exit
...
-----
A:ALA-41>config>service#
```

2.3.3.7 Configuring VRRP

Configuring VRRP parameters on an IES interface is optional and applies only to the 7450 ESS and 7750 SR. VRRP can be configured in either an owner or non-owner mode. The owner is the VRRP router whose virtual router IP address is the same as the real interface IP address. This is the router that responds to packets addressed to one of the IP addresses for ICMP pings, TCP connections, etc. All other virtual router instances participating in this message domain must have the same VRID configured and cannot be configured as owner.

For further information about VRRP CLI syntax and command descriptions refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide*.

The following example displays the IES configuration:

```
*A:ALA-A>config>service>ies>if# info
-----
          address 10.10.36.2/24
          vrrp 2 owner
              backup 10.10.36.2
          authentication-key "3WErEDozxyQ" hash
          exit
-----
*A:ALA-A>config>service#
```

2.3.3.8 Configuring IPSec Parameters

The following output displays an IES service with IPSec parameters configured for the 7750 SR:

```
*A:ALA-49>config# info
-----
...
    service
        ies 100 customer 1 create
            interface "ipsec-public" create
                address 10.10.10.1/24
                sap ipsec-1.public:1 create
                exit
            exit
        no shutdown
    exit
exit
...
-----
*A:ALA-49>config#
```

2.3.3.9 IGMP Host Tracking

The following output displays an IES service with IGMP host tracking parameters configured.

```
*A:ALA-49>config>service# info
-----
...
    ies 25 customer 1 create
        interface "ip_if_4" create
            loopback
            hold-time down ip 1200
            address 10.64.64.64/24
            sap lag-64:64 create
                no shutdown
            exit
            allow-directed-broadcasts
            host-connectivity-verify
            ip-mtu 9000
            local-dhcp-server "server 1"
            local-proxy-arp
            proxy-arp-policy treetrace-1
            remote-proxy-arp
            secondary 10.3.4.5 255.255.255.0
            secondary 10.3.4.5/24
            tos-marking-state trusted
            tos-marking-state untrusted
            urpf-check
            exit
        exit
        igmp-host-tracking
            expiry-time 65535
            no shutdown
        exit
    ...
-----
*A:ALA-49>config>service#
```

2.4 Service Management Tasks

This section describes IES service management tasks:

2.4.1 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 parameter(s) (such as description, SAP information and SDP information) and then enter the new information.

The following example displays the modified service for 7450 ESS and 7750 SR:

```
*A:ALA-A>config>service>ies# info
-----
ies 1000 customer 50 vpn 1000 create
    description "This is a new description"
    interface "to-web" create
        address 10.1.1.1/24
        mac 00:dc:98:1d:00:00
        allow-directed-broadcast
        sap 22/1/50:0 create
    exit
    exit
    no shutdown
exit
-----
*A:ALA-A>config>service#
```

The following example displays the modified service for the 7950 XRS:

```
*A:ALA-A>config>service>ies# info
-----
ies 1000 customer 50 vpn 1000 create
    description "This is a new description"
    interface "to-web" create
        address 10.1.1.1/24
        mac 00:dc:98:1d:00:00
        allow-directed-broadcast
    exit
    exit
    no shutdown
exit
-----
*A:ALA-A>config>service#
```

2.4.2 Deleting a Spoke-SDP

To delete the spoke SDP from the service interface must be shut down. This cleans up the associated VC labels.

Use the following CLI syntax to delete a spoke SDP from an interface for the 7450 ESS and 7750 SR:

CLI Syntax:

```
config>service# ies service-id [customer customer-id]
[vpn vpn-id]
interface ip-int-name
[no] spoke-sdp sdp-id:vc-id
shutdown
```

The following example displays the spoke SDP configuration for the 7450 ESS and 7750 SR:

```
A:ALA-48>config>service>ies# info
-----
description "to internet"
interface "spokeSDP-test" create
exit
no shutdown
-----
A:ALA-48>config>service>ies#
```

2.4.3 Deleting an IES Service

An IES service cannot be deleted until SAPs and interfaces are shut down *and* deleted and the service is shutdown on the service level.

Use the following CLI syntax to delete an IES service:

CLI Syntax:

```
config>service#
[no] ies service-id
shutdown
[no] interface ip-int-name
shutdown
[no] sap sap-id
shutdown
```


2.4.4 Disabling an IES Service

An IES service can be shut down without deleting the service parameters.

CLI Syntax: `config>service> ies service-id
shutdown`

2.4.5 Re-Enabling an IES Service

To re-enable an IES service that was shut down.

CLI Syntax: `config>service> ies service-id
[no] shutdown`

Example: `config>service# ies 2000
config>service>ies# no shutdown
config>service>ies# exit`

2.5 IES Services Command Reference

This section describes the IES services command reference.

2.5.1 Command Hierarchies

- [Global Commands](#)
- [AARP Interface Commands](#)
- [IES Service Interface Commands](#)
- [Routed VPLS Commands](#)
- [Redundant Interface Commands](#)
- [Interface SAP Commands](#)
- [VRRP Commands](#)
- [Spoke SDP Commands](#)
- [Subscriber Interface Commands](#)
 - [Group Interface Commands](#)
 - [Group Interface SAP Commands](#)
 - [Group Interface SAP ETH-CFM Commands](#)

2.5.1.1 Global Commands

```
config
  — service
    — ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
    — no ies service-id
      — description description-string
      — no description [description-string]
      — igmp-host-tracking
        — expiry-time expiry-time
        — no expiry-time
        — [no] shutdown
      — [no] shutdown
```

2.5.1.2 AARP Interface Commands

```
config
```

```

— service
  — ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
  — no ies service-id
    — aarp-interface arp-int-name [create]
    — no aarp-interface arp-int-name
      — description long-description-string
      — no description
      — ip-mtu octets
      — no ip-mtu
      — [no] shutdown
      — spoke-sdp sdp-id:vc-id [create]
      — no spoke-sdp sdp-id:vc-id
        — aarp aarp-id type {subscriber-side-shunt | network-side-shunt}
        — no aarp
        — description description-string
        — no description
        — egress
          — filter ip ip-filter-id
          — no filter
          — vc-label vc-label
          — no vc-label [vc-label]
        — ingress
          — filter ip ip-filter-id
          — no filter
          — vc-label vc-label
          — no vc-label [vc-label]
        — shutdown

```

2.5.1.3 IES Service Interface Commands

```

config
  — service
    — ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
    — interface ip-int-name [create]
    — interface ip-int-name [create] tunnel
    — no interface ip-int-name
      — address {ip-address/mask | ip-address netmask} [broadcast {all-ones | host-ones}] [track-srrp srrp-instance]
      — no address [ip-address/mask | ip-address netmask]
      — [no] allow-directed-broadcasts
      — [no] arp-learn-unsolicited
      — arp-limit limit [log-only] [threshold percent]
      — no arp-limit
      — [no] arp-populate
      — [no] arp-populate-host-route
      — [no] arp-proactive-refresh
      — arp-retry-timer timer-multiple
      — no arp-retry-timer
      — arp-route-tag tag
      — no arp-route-tag

```

-
- **arp-timeout** *seconds*
 - **no arp-timeout** [*seconds*]
 - **authentication-policy** *name*
 - **no authentication-policy**
 - **bfd** *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier*]
[**echo-receive** *echo-interval*] [**type** *cpm-np*]
 - **no bfd**
 - **cflowd-parameters**
 - **sampling** {*unicast* | *multicast*} **type** {*acl* | *interface*} [**direction** {*ingress-only* | *egress-only* | *both*}]
 - **no sampling** {*unicast* | *multicast*}
 - **cpu-protection** *policy-id*
 - **no cpu-protection**
 - **description** *long-description-string*
 - **no description** [*long-description-string*]
 - **dhcp**
 - **description** *description-string*
 - **no description**
 - **gi-address** *ip-address* [*src-ip-addr*]
 - **no gi-address**
 - **lease-populate** [*nbr-of-leases*]
 - **no lease-populate**
 - **[no] option**
 - **action** {*replace* | *drop* | *keep*}
 - **no action**
 - **circuit-id** [*ascii-tuple* | *ifindex* | *sap-id* | *vlan-ascii-tuple*]
 - **no circuit-id**
 - **remote-id** [*mac* | *string* *string*]
 - **no remote-id**
 - **[no] vendor-specific-option**
 - **[no] client-mac-address**
 - **[no] sap-id**
 - **[no] service-id**
 - **string** *text*
 - **no string**
 - **[no] system-id**
 - **proxy-server**
 - **emulated-server** *ip-address*
 - **no emulated-server**
 - **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*] [**override**]
 - **no lease-time**
 - **[no] shutdown**
 - **python-policy** *policy-name*
 - **no python-policy**
 - **relay-proxy** [**release-update-src-ip**] [**siaddr-override** *ip-address*]
 - **no relay-proxy**
 - **server** *server1* [*server2*]
 - **no server**
 - **[no] shutdown**
 - **[no] trusted**
 - **dynamic-tunnel-redundant-next-hop** *ip-address*

-
- **no dynamic-tunnel-redundant-next-hop**
 - **[no] enable-ingress-stats**
 - **[no] enable-mac-accounting**
 - **hold-time**
 - **down ip seconds** [init-only]
 - **no down ip**
 - **down ipv6 seconds** [init-only]
 - **no down ipv6**
 - **up ip seconds**
 - **no up ip**
 - **up ipv6 seconds**
 - **no up ipv6**
 - **host-connectivity-verify** [source {vrrp | interface}] [interval interval] [action {remove | alarm} [timeout retry-timeout] [retry-count count]]
 - **icmp**
 - **[no] mask-reply**
 - **param-problem** [number seconds]
 - **no param-problem**
 - **redirects** number seconds
 - **no redirects** [number seconds]
 - **ttl-expired** number seconds
 - **no ttl-expired** [number seconds]
 - **unreachables** number seconds
 - **no unreachables** [number seconds]
 - **if-attribute**
 - **[no] admin-group** group-name [group-name]
 - **no admin-group**
 - **[no] srlg-group** group-name [group-name]
 - **no srlg-group**
 - **ingress**
 - **policy-accounting** template-name
 - **no policy-accounting**
 - **ip-helper-address** gateway-address
 - **no ip-helper-address**
 - **ip-mtu** octets
 - **no ip-mtu**
 - **ipcp**
 - **dns ip-address** [secondary ip-address]
 - **dns secondary ip-address**
 - **no dns** [ip-address] [secondary ip-address]
 - **peer-ip-address** ip-address
 - **no peer-ip-address**
 - **[no] ipv6**
 - **address** ipv6-address/prefix-length [eui-64] [track-srrp srrp-instance] [modifier cga-modifier] [dad-disable] [primary-preference primary-preference]
 - **no address** ipv6-address/prefix-length
 - **bfd** transmit-interval [receive receive-interval] [multiplier multiplier] [echo-receive echo-interval] [type cpm-np]
 - **no bfd**
 - **[no] dad-disable**
 - **[no] dhcp6-relay**
 - **description** description-string

-
- **no description**
 - **lease-populate** *[nbr-of-leases]*
 - **lease-populate** *[nbr-of-leases]* **route-populate** *[pd]* **na**
[ta]
 - **lease-populate** *[nbr-of-leases]* **route-populate** **pd** *[na]*
[ta] **[exclude]**
 - **lease-populate** *[nbr-of-leases]* **route-populate** *[pd]* *[na]*
ta
 - **no lease-populate**
 - **link-address** *ipv6-address*
 - **no link-address**
 - **[no] neighbor-resolution**
 - **[no] option**
 - **interface-id** *[ascii-tuple]*
 - **interface-id** *ifindex*
 - **interface-id** *sap-id*
 - **interface-id** *string* *string*
 - **no interface-id**
 - **[no] remote-id**
 - **python-policy** *policy-name*
 - **no python-policy**
 - **server** *ipv6-address* *[ipv6-address]*
 - **no server** *[ipv6-address]*
 - **[no] shutdown**
 - **source-address** *ipv6-address*
 - **no source-address**
 - **[no] dhcp6-server**
 - **max-nbr-of-leases** *max-nbr-of-leases*
 - **no max-nbr-of-leases**
 - **[no] prefix-delegation**
 - **[no] prefix** *ipv6-address/prefix-length*
 - **duid** *duid* *[iaid iaaid]*
 - **no duid**
 - **preferred-lifetime** *seconds*
 - **preferred-lifetime** **infinite**
 - **no preferred-lifetime**
 - **valid-lifetime** *seconds*
 - **valid-lifetime** **infinite**
 - **no valid-lifetime**
 - **[no] shutdown**
 - **icmp6**
 - **packet-too-big** *[number seconds]*
 - **no packet-too-big**
 - **param-problem** *[number seconds]*
 - **no param-problem**
 - **redirects** *[number seconds]*
 - **no redirects**
 - **time-exceeded** *[number seconds]*
 - **no time-exceeded**
 - **unreachables** *[number seconds]*
 - **no unreachables**
 - **link-local-address** *ipv6-address* **[dad-disable]**
 - **no link-local-address**
 - **local-dhcp-server** *local-server-name*

```

— no local-dhcp-server
— [no] local-proxy-nd
— nd-learn-unsolicited {global | link-local | both}
— no nd-learn-unsolicited
— [no] nd-populate-host-route
— nd-proactive-refresh {global | link-local | both}
— no nd-proactive-refresh
— nd-route-tag tag
— no nd-route-tag
— neighbor ipv6-address mac-address
— no neighbor ipv6-address
— neighbor-limit limit [log-only] [threshold percent]
— no neighbor-limit
— proxy-nd-policy policy-name [policy-name]
— no proxy-nd-policy
— qos-route-lookup [source | destination]
— no qos-route-lookup
— [no] secure-nd
    — [no] allow-unsecured-msgs
    — link-local-modifier modifier
    — no link-local-modifier
    — public-key-min-bits bits
    — no public-key-min-bits
    — security-parameter sec
    — no security-parameter
    — [no] shutdown
— stale-time seconds
— no stale-time
— tcp-mss mss-value
— no tcp-mss
— [no] urpf-check
    — mode {strict | loose | strict-no-ecmp}
— load-balancing
    — egr-ip-load-balancing {source | destination | inner-ip}
    — no egr-ip-load-balancing
    — [no] spi-load-balancing
    — [no] teid-load-balancing
— [no] local-dhcp-server
— [no] local-proxy-arp
— [no] loopback
— mac ieee-address
— no mac [ieee-address]
— monitor-oper-group name
— no monitor-oper-group
— multicast-network-domain multicast-network-domain
— no multicast-network-domain
— proxy-arp-policy policy-name [policy-name]
— no proxy-arp-policy
— [no] ptp-hw-assist
— qos-route-lookup [source | destination]
— no qos-route-lookup
— [no] remote-proxy-arp
— secondary {ip-address/mask | ip-address netmask} [broadcast {all-
    ones | host-ones}] [igp-inhibit] [track-srrp srrp-instance]

```


- **no secondary** {*ip-address/mask* | *ip-address netmask*}
- **shcv-policy-ipv4** *policy-name*
- **no shcv-policy-ipv4**
- **shcv-policy-ipv6** *policy-name*
- **no shcv-policy-ipv6**
- **[no] shutdown**
- **static-arp** *ieee-mac-addr* **unnumbered**
- **static-arp** *ip-address* *ieee-mac-address*
- **no static-arp** [*ieee-mac-addr*] **unnumbered**
- **no static-arp** *ip-address* [*ieee-mac-address*]
- **static-tunnel-redundant-next-hop** *ip-address*
- **no static-tunnel-redundant-next-hop**
- **tcp-mss** *mss-value*
- **no tcp-mss**
- **tos-marking-state** {*trusted* | *untrusted*}
- **no tos-marking-state**
- **unnumbered** [{*ip-int-name* | *ip-address*}]
- **no unnumbered**
- **[no] urpf-check**
 - **mode** {*strict* | *loose* | *strict-no-ecmp*}
- **vas-if-type** {*to-from-access* | *to-from-network* | *to-from-both*}
- **no vas-if-type**

2.5.1.4 Routed VPLS Commands

- ```

config
 — service
 — ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
 — interface ip-int-name [create]
 — interface ip-int-name [create] tunnel
 — no interface ip-int-name
 — vpls service-name
 — no vpls
 — egress
 — reclassify-using-qos policy-id
 — no reclassify-using-qos
 — v4-routed-override-filter ipv4-filter-id
 — no v4-routed-override-filter
 — v6-routed-override-filter ipv6-filter-id
 — no v6-routed-override-filter
 — ingress
 — v4-routed-override-filter ipv4-filter-id
 — no v4-routed-override-filter
 — v6-routed-override-filter ipv6-filter-id
 — no v6-routed-override-filter

```

### 2.5.1.5 Redundant Interface Commands

```

config
 — service
 — ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
 — redundant-interface ip-int-name [create]
 — no redundant-interface ip-int-name
 — address {ip-address/mask | ip-address netmask} [remote-ip ip-address]
 — no address
 — description long-description-string
 — no description [long-description-string]
 — hold-time
 — down ip seconds [init-only]
 — no down ip
 — down ipv6 seconds [init-only]
 — no down ipv6
 — up ip seconds
 — no up ip
 — up ipv6 seconds
 — no up ipv6
 — [no] shutdown
 — spoke-sdp sdp-id:vc-id [create]
 — no spoke-sdp sdp-id:vc-id
 — egress
 — filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
 — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
 — vc-label egress-vc-label
 — no vc-label [egress-vc-label]
 — ingress
 — filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
 — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
 — vc-label ingress-vc-label
 — no vc-label [ingress-vc-label]

```

### 2.5.1.6 Interface SAP Commands

```

config
 — service
 — ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
 — interface ip-int-name [create]
 — interface ip-int-name [create] tunnel
 — no interface ip-int-name
 — sap sap-id [create]
 — no sap sap-id
 — aarp aarpId type type
 — no aarp
 — accounting-policy acct-policy-id
 — no accounting-policy [acct-policy-id]
 — anti-spoof {ip | ip-mac | mac}

```

- **no anti-spoof**
- **app-profile** *app-profile-name*
- **no app-profile**
- **atm**
  - **egress**
    - **traffic-desc** *traffic-desc-profile-id*
    - **no traffic-desc**
  - **encapsulation** *atm-encap-type*
  - **ingress**
    - **traffic-desc** *traffic-desc-profile-id*
    - **no traffic-desc**
  - **oam**
    - [no] **alarm-cells**
    - [no] **periodic-loopback**
- **bandwidth** *bandwidth*
- **no bandwidth**
- **calling-station-id** *calling-station-id*
- **no calling-station-id**
- [no] **collect-stats**
- **cpu-protection** *policy-id* {[**mac-monitoring**] | [**eth-cfm-monitoring**] [**aggregate**] [**car**] | [**ip-src-monitoring**]}
- **no cpu-protection**
- **description** *long-description-string*
- **no description** [*long-description-string*]
- **dist-cpu-protection** *policy-name*
- **no dist-cpu-protection**
- **egress**
  - [no] **agg-rate**
    - [no] **limit-unused-bandwidth**
    - [no] **queue-frame-based-accounting**
    - **rate** *rate*
    - **no rate**
  - **filter ip** *ip-filter-id*
  - **filter ipv6** *ipv6-filter-id*
  - **no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
  - [no] **hsmdda-queue-override**
    - **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
    - **no packet-byte-offset**
    - **queue** *queue-id* [**create**]
    - **no queue** *queue-id*
      - **mbs** {*size* [**kilobytes**] | **bytes** | **default**}
      - **no mbs**
      - **rate** *pir-rate*
      - **no rate**
      - **slope-policy** *hsmdda-slope-policy-name*
      - **no slope-policy**
      - **wrr-weight** *weight*
      - **no wrr-weight**
    - **secondary-shaper** *secondary-shaper-name*
    - **no secondary-shaper**
    - **wrr-policy** *hsmdda-wrr-policy-name*
  - [no] **policer-override**
    - **policer** *policer-id* [**create**]

- 
- **no policer** *policer-id*
  - **[no] qinq-mark-top-only**
  - **qos** *policy-id* [**port-redirect-group** *queue-group-name* **instance** *instance-id*]
  - **no qos** [*policy-id*]
  - **queue-group-redirect-list** *redirect-list-name*
  - **no queue-group-redirect-list**
  - **[no] queue-override**
    - **hs-secondary-shaper** *policy-name*
    - **no hs-secondary-shaper**
    - **hs-wrr-group** *group-id* [**create**]
    - **no hs-wrr-group** *group-id*
      - **class-weight** *weight*
      - **no class-weight**
      - **percent-rate** *percent*
      - **no percent-rate**
      - **rate** *rate*
      - **no rate**
    - **queue** *queue-id* [**create**]
    - **no queue** *queue-id*
      - **adaptation-rule** [**pir** {**max** | **min** | **closest**}] [**cir** {**max** | **min** | **closest**}]
      - **avg-frame-overhead** *percentage*
      - **no avg-frame-overhead**
      - **burst-limit** {**default** | **size** [**bytes** | **kilobytes**]}
      - **no burst-limit**
      - **cbs** *size-in-kbytes*
      - **no cbs**
      - **drop-tail low**
        - **percent-reduction-from-mbs** *percent*
        - **no percent-reduction-from-mbs**
      - **hs-class-weight** *weight*
      - **no hs-class-weight**
      - **hs-wred-queue** *policy* *slope-policy-name*
      - **no hs-wred-queue**
      - **hs-wrr-weight** *weight*
      - **no hs-wrr-weight**
      - **mbs** *size* {[**bytes** | **kilobytes**] | **default**}
      - **no mbs**
      - **[no] monitor-depth**
      - **parent** [**weight** *weight*] [**cir-weight** *cir-weight*]
      - **percent-rate** *pir-percent* [**cir** *cir-percent*]
      - **no percent-rate**
      - **rate** *pir-rate* [**cir** *cir-rate*]
      - **no rate**
    - **[no] scheduler-override**
      - **scheduler** *scheduler-name* [**create**]
      - **no scheduler** *scheduler-name*
        - **parent** [**weight** *weight*] [**cir-weight** *cir-weight*]
        - **no parent**

- **rate** *pir-rate* [**cir** *cir-rate*]
  - **no rate**
  - **scheduler-policy** *scheduler-policy-name*
  - **no scheduler-policy**
- **eth-cfm**
  - **collect-lmm-fc-stats**
  - **[no] collect-lmm-stats**
    - **fc** *fc-name* [*fc-name*]
    - **no fc**
    - **fc-in-profile** *fc-name* [*fc-name*]
    - **no fc-in-profile**
  - **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**
      - **[no] interface-support-enable**
    - **[no] ccm-enable**
    - **ccm-ltm-priority** *priority*
    - **no ccm-ltm-priority**
    - **ccm-padding-size** *ccm-padding*
    - **no ccm-padding-size**
    - **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**
    - **fault-propagation-enable** {**use-if-tlv** | **suspend-ccm**}
    - **no fault-propagation-enable**
    - **grace**
      - **eth-ed**
        - **max-rx-defect-window** *seconds*
        - **no max-rx-defect-window**
        - **priority** *priority*
        - **no priority**
        - **[no] rx-eth-ed**
        - **[no] tx-eth-ed**
      - **eth-vsm-grace**
        - **[no] rx-eth-vsm-grace**
        - **[no] tx-eth-vsm-grace**
    - **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}
    - **one-way-delay-threshold** *seconds*
    - **[no] shutdown**
  - **snelch-ingress-levels** [*md-level*] [*md-level*]
  - **no snelch-ingress-levels**
  - **tunnel-fault** [**accept** | **ignore**]
- **frame-relay**
  - **[no] frf-12**
    - **ete-fragment-threshold** *fragment-threshold*
    - **no ete-fragment-threshold**
    - **[no] interleave**

- 
- **scheduling-class** *class-id*
  - **no scheduling-class**
  - **fwd-wholesale**
    - **pppoe** *service-id*
    - **no pppoe**
  - **host-lockout-policy** *policy-name*
  - **no host-lockout-policy**
  - **[no] host-shutdown**
  - **ingress**
    - **filter ip** *ip-filter-id*
    - **filter ipv6** *ipv6-filter-id*
    - **no filter** [*ip ip-filter-id*] [*ipv6 ipv6-filter-id*]
    - **match-qinq-dot1p** {*top* | *bottom*}
    - **no match-qinq-dot1p**
    - **[no] policer-override**
      - **policer** *policer-id* [*create*]
      - **no policer** *policer-id*
    - **qos** *policy-id* [*shared-queuing* | *multipoint-shared*] [*fp-redirect-group queue-group-name instance instance-id*]
    - **no qos** [*policy-id*]
    - **queue-group-redirect-list** *redirect-list-name*
    - **no queue-group-redirect-list**
    - **[no] queue-override**
      - **queue** *queue-id* [*create*]
      - **no queue** *queue-id*
        - **adaptation-rule** [*pir* {*max* | *min* | *closest*}] [*cir* {*max* | *min* | *closest*}]
        - **no adaptation-rule**
        - **avg-frame-overhead** *percentage*
        - **no avg-frame-overhead**
        - **cbs** *size-in-kbytes*
        - **no cbs**
        - **drop-tail low**
          - **percent-reduction-from-mbs** *percent*
          - **no percent-reduction-from-mbs**
        - **mbs** *size* {[*bytes* | *kilobytes*] | *default*}
        - **no mbs**
        - **[no] monitor-depth**
        - **parent** [*weight weight*] [*cir-weight cir-weight*]
        - **percent-rate** *pir-percent* [*cir cir-percent*]
        - **no percent-rate**
        - **rate** *pir-rate* [*cir cir-rate*]
        - **no rate**
    - **[no] scheduler-override**
      - **[no] scheduler** *scheduler-name*
        - **parent** [*weight weight*] [*cir-weight cir-weight*]
        - **no parent**
        - **rate** *pir-rate* [*cir cir-rate*]
        - **no rate**
    - **scheduler-policy** *scheduler-policy-name*

- 
- **no scheduler-policy**
  - **ip-tunnel** *name* [**create**]
  - **no ip-tunnel** *name*
    - **backup-remote-ip** *ip-address*
    - **no backup-remote-ip**
    - **[no] clear-df-bit**
    - **delivery-service** *service-id*
    - **delivery-service** *name* *service-name*
    - **no delivery-service**
    - **description** *description-string*
    - **no description**
    - **[no] dest-ip** *ip-address*
    - **dscp** *dscp-name*
    - **no dscp**
    - **gre-header**
    - **gre-header** *send-key* *send-key* *receive-key* *receive-key*
    - **no gre-header**
    - **ip-mtu** *octets*
    - **no ip-mtu**
    - **reassemble** [*wait-msecs*]
    - **no reassemble**
    - **remote-ip** *ip-address*
    - **no remote-ip**
    - **[no] shutdown**
    - **source** *ip-address*
    - **no source**
  - **ipsec-gw** *name*
  - **no ipsec-gw**
    - **cert**
      - **cert-profile** *name*
      - **no cert-profile**
      - **status-verify**
        - **default-result** {*revoked* | *good*}
        - **no default-result**
        - **primary** *primary* **secondary** *secondary*
      - **trust-anchor-profile** *file-name*
      - **no trust-anchor-profile**
    - **default-secure-service** *service-id* **interface** *ip-int-name*
    - **default-secure-service** *name* *service-name* **interface** *ip-int-name*
    - **no default-secure-service**
    - **default-tunnel-template** *ipsec template identifier*
    - **no default-tunnel-template**
    - **local-gateway-address** *ip-address*
    - **no local-gateway-address**
    - **local-id** **type** {*ipv4* *v4address* | *fqdn* *fqdn-value* | *ipv6* *v6address*}
    - **no local-id**
    - **pre-shared-key** *key* [*hash* | *hash2*]
    - **no pre-shared-key**
    - **[no] shutdown**
  - **lag-link-map-profile** *lag-link-map-profile-id*
  - **no lag-link-map-profile**
  - **lag-per-link-hash** **class** {*1* | *2* | *3*} **weight** *weight*

- **no lag-per-link-hash**
- **multi-service-site** *customer-site-name*
- **no multi-service-site**
- **[no] shutdown**
- **static-host** **ip** *ip/did-address* [**mac** *ieee-address*] [**create**]
- **static-host** **mac** *ieee-address* [**create**]
- **no static-host** [**ip** *ip-address*] **mac** *ieee-address*
- **no static-host** **all** [**force**]
- **no static-host** **ip** *ip-address*
  - **ancp-string** *ancp-string*
  - **no ancp-string**
  - **app-profile** *app-profile-name*
  - **no app-profile**
  - **inter-dest-id** *intermediate-destination-id*
  - **no inter-dest-id**
  - **[no] shutdown**
  - **sla-profile** *sla-profile-name*
  - **no sla-profile**
  - **sub-profile** *sub-profile-name*
  - **no sub-profile**
  - **subscriber** *sub-ident*
  - **no subscriber**
  - **[no] subscriber-sap-id**
- **transit-policy** {**ip** *ip-aasub-policy-id* | **prefix** *prefix-aasub-policy-id*}
- **no transit-policy**

## 2.5.1.7 VRRP Commands

- ```

config
  — service
    — ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
      — interface ip-int-name [create]
      — interface ip-int-name [create] tunnel
      — no interface ip-int-name
        — [no] ipv6
          — vrrp virtual-router-id [owner] [passive]
          — no vrrp virtual-router-id
            — [no] backup ip-address
            — [no] bfd-enable interface interface-name dst-ip ip-address
            — [no] bfd-enable service-id interface interface-name dst-ip ip-address
            — [no] bfd-enable interface interface-name dst-ip ip-address name name
            — init-delay seconds
            — no init-delay
            — mac mac-address
            — no mac
            — [no] master-int-inherit

```

```

— message-interval {[seconds] [milliseconds milliseconds]}
— no message-interval
— oper-group group-name
— no oper-group
— [no] ping-reply
— policy vrrp-policy-id
— no policy
— [no] preempt
— priority base-priority
— no priority
— [no] shutdown
— [no] standby-forwarding
— [no] telnet-reply
— [no] traceroute-reply
— vrrp virtual-router-id [owner] [passive]
— no vrrp virtual-router-id
— authentication-key {authentication-key | hash-key} [hash | hash2]
— no authentication-key
— [no] backup ip-address
— [no] bfd-enable interface interface-name dst-ip ip-address
— [no] bfd-enable service-id interface interface-name dst-ip ip-address address
— [no] bfd-enable interface interface-name dst-ip ip-address name name
— init-delay seconds
— no init-delay
— mac ieee-address
— no mac
— [no] master-int-inherit
— message-interval {[seconds] [milliseconds milliseconds]}
— no message-interval
— oper-group group-name
— no oper-group
— [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

```

2.5.1.8 Spoke SDP Commands

```

config
— service

```

-
- **ies** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**name** *name*]
 - **interface** *ip-int-name* [**create**]
 - **interface** *ip-int-name* [**create**] **tunnel**
 - **no interface** *ip-int-name*
 - **spoke-sdp** *sdp-id:vc-id* [**vc-type** {*ether* | *ipipe*}] [**create**]
 - **no spoke-sdp** *sdp-id:vc-id*
 - **aarp** *aarpId type* *type*
 - **no aarp**
 - **accounting-policy** *acct-policy-id*
 - **no accounting-policy**
 - **app-profile** *app-profile-name*
 - **no app-profile**
 - [no] **bfd-enable**
 - **bfd-template** *name*
 - **no bfd-template**
 - [no] **collect-stats**
 - **control-channel-status**
 - [no] **acknowledgment**
 - **refresh-timer** *value*
 - **no refresh-timer**
 - **request-timer** *timer1* **retry-timer** *timer2* **timeout-**
multiplier *multiplier*
 - **no request-timer**
 - [no] **shutdown**
 - [no] **control-word**
 - **description** *long-description-string*
 - **no description**
 - **egress**
 - **filter** **ip** *ip-filter-id*
 - **filter** **ipv6** *ipv6-filter-id*
 - **no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
 - **qos** *network-policy-id* **port-redirect-group** *queue-group-*
name [**instance** *instance-id*]
 - **no qos** [*network-policy-id*]
 - **vc-label** *egress-vc-label*
 - **no vc-label** [*egress-vc-label*]
 - [no] **entropy-label**
 - **eth-cfm**
 - **collect-lmm-fc-stats**
 - **fc** *fc-name* [*fc-name*]
 - **no fc**
 - **fc-in-profile** *fc-name* [*fc-name*]
 - **no fc-in-profile**
 - [no] **collect-lmm-stats**
 - **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**
 - [no] **interface-support-enable**
 - [no] **ccm-enable**
 - **ccm-ltm-priority** *priority*
 - **no ccm-ltm-priority**
 - **ccm-padding-size** *ccm-padding*
 - **no ccm-padding-size**

- **description** *long-description-string*
- **no description**
- **[no] eth-test-enable**
 - **bit-error-threshold** *bit-errors*
 - **test-pattern** {*all-zeros* | *all-ones*} [*crc-enable*]
 - **no test-pattern**
- **fault-propagation-enable** {*use-if-tlv* | *suspend-ccm*}
- **no fault-propagation-enable**
- **grace**
 - **eth-ed**
 - **max-rx-defect-window** *seconds*
 - **no max-rx-defect-window**
 - **priority** *priority*
 - **no priority**
 - **[no] rx-eth-ed**
 - **[no] tx-eth-ed**
 - **eth-vsm-grace**
 - **[no] rx-eth-vsm-grace**
 - **[no] tx-eth-vsm-grace**
- **low-priority-defect** {*allDef* | *macRemErrXcon* | *remErrXcon* | *errXcon* | *xcon* | *noXcon*}
- **low-priority-defect** *seconds*
- **one-way-delay-threshold** *seconds*
- **[no] shutdown**
- **squelch-ingress-levels** [*md-level*] [*md-level*]
- **no squelch-ingress-levels**
- **hash-label** [*signal-capability*]
- **no hash-label**
- **ingress**
 - **filter ip** *ip-filter-id*
 - **filter ipv6** *ipv6-filter-id*
 - **no filter** [*ip ip-filter-id*] [*ipv6 ipv6-filter-id*]
 - **qos** *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*
 - **no qos** [*network-policy-id*]
 - **vc-label** *ingress-vc-label*
 - **no vc-label** [*ingress-vc-label*]
- **[no] pw-path-id**
 - **agi** *agi*
 - **no agi**
 - **saii-type2** *global-id:node-id:ac-id*
 - **no saii-type2**
 - **taii-type2** *global-id:node-id:ac-id*
 - **no taii-type2**
- **[no] shutdown**
- **transit-policy** {*ip ip-aasub-policy-id* | **prefix** *prefix-aasub-policy-idb-policy-id*}
- **no transit-policy**

2.5.1.9 Subscriber Interface Commands

```

config
  — service
    — ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
      — [no] subscriber-interface ip-int-name
        — address ip-address/mask [netmask] [gw-ip-address ip-address]
          [populate-host-routes] [track-srrp srrp-instance] [holdup-time
            milli-seconds]
        — no address ip-address/mask [netmask]
        — [no] allow-unmatching-subnets
        — description long-description-string
        — no description [long-description-string]
        — dhcp
          — client-applications {[dhcp] [ppp]}
          — no client-applications
          — description description-string
          — no description
          — gi-address ip-address [src-ip-addr]
          — no gi-address
          — python-policy policy-name
          — no python-policy
          — relay-proxy [release-update-src-ip] [siaddr-override ip-
            address]
          — no relay-proxy
          — server server1 [server2]
          — no server
          — [no] shutdown
        — hold-time
          — up ip seconds
          — no up ip
          — up ipv6 seconds
          — no up ipv6
          — down ip seconds [init-only]
          — no down ip
          — down ipv6 seconds [init-only]
          — no down ipv6
        — [no] ipv6
          — [no] allow-unmatching-prefixes
          — delegated-prefix-length prefix-length
          — delegated-prefix-length variable
          — no delegated-prefix-length
          — [no] subscriber-prefixes
            — prefix ipv6-address/prefix-length [pd | wan-host]
            — no prefix ipv6-address/prefix-length
        — [no] unnumbered {ip-address | inf-name}
          — [no] wpp
            — initial-app-profile profile-name
            — no initial-app-profile
            — initial-sla-profile profile-name
            — no initial-sla-profile
            — initial-sub-profile profile-name
            — no initial-sub-profile

```

- **portal** *router router-instance name wpp-portal-name*
- **no portal**
- **[no] restore-disconnected**
- **[no] shutdown**

2.5.1.9.1 Group Interface Commands

```

config
— service
— ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
— [no] subscriber-interface ip-int-name
— group-interface ip-int-name [create]
— group-interface ip-int-name [create] lns
— group-interface ip-int-name [create] wlangw
— no group-interface ip-int-name
— arp-host
— host-limit max-num-hosts
— no host-limit
— min-auth-interval min-auth-interval
— no min-auth-interval
— sap-host-limit max-num-hosts-sap
— no sap-host-limit
— [no] shutdown
— [no] arp-populate
— arp-timeout seconds
— no arp-timeout [seconds]
— authentication-policy name
— no authentication-policy
— cflowd-parameters
— sampling {unicast | multicast} type {acl | interface}
— [direction {ingress-only | egress-only | both}]
— no sampling {unicast | multicast}
— description long-description-string
— no description [long-description-string]
— dhcp
— client-applications {[dhcp] [ppp]}
— no client-applications
— description description-string
— no description
— [no] enable-ingress-stats
— filter filter-id
— no filter
— gi-address ip-address [src-ip-addr]
— no gi-address
— lease-populate nbr-of-leases
— no lease-populate
— [no] match-circuit-id
— [no] option
— action {replace | drop | keep}
— no action

```

-
- **circuit-id** [ascii-tuple | ifindex | sap-id | vlan-ascii-tuple]
 - **no circuit-id**
 - **remote-id** [mac | string *string*]
 - **no remote-id**
 - [no] **vendor-specific-option**
 - [no] **client-mac-address**
 - [no] **pool-name**
 - [no] **sap-id**
 - [no] **service-id**
 - **string** *text*
 - **no string**
 - [no] **system-id**
 - **proxy-server**
 - **emulated-server** *ip-address*
 - **no emulated-server**
 - **lease-time** [days *days*] [hrs *hours*] [min *minutes*] [sec *seconds*] [override]
 - **no lease-time**
 - [no] **shutdown**
 - **relay-proxy** [release-update-src-ip] [siaddr-override *ip-address*]
 - **no relay-proxy**
 - **server** *server1* [*server2*]
 - **no server**
 - [no] **shutdown**
 - [no] **trusted**
 - **user-db** *local-user-db-name*
 - **no user-db**
 - **host-connectivity-verify** [interval *interval*] [action {remove | alarm}] [timeout *retry-timeout*] [retry-count *count*] [family *family*]
 - **icmp**
 - [no] **mask-reply**
 - **redirects** [*number seconds*]
 - **no redirects**
 - **ttl-expired** [*number seconds*]
 - **no ttl-expired**
 - **unreachables** [*number seconds*]
 - **no unreachables**
 - **ip-mtu** *octets*
 - **no ip-mtu**
 - [no] **ipv6**
 - [no] **allow-unmatching-prefixes**
 - **delegated-prefix-length** *bits*
 - **delegated-prefix-length** **variable**
 - **no delegated-prefix-length**
 - [no] **router-advertisements**
 - **current-hop-limit** *hop-count*
 - **no current-hop-limit**
 - [no] **managed-configuration**
 - **max-advertisement** *seconds*
 - **no max-advertisement**
 - **min-advertisement** *seconds*

-
- **no min-advertisement**
 - **mtu** *bytes*
 - **no mtu**
 - [no] **other-stateful-configuration**
 - [no] **prefix-options**
 - [no] **autonomous**
 - **preferred-lifetime** [*seconds* | *infinite*]
 - **no preferred-lifetime**
 - **valid-lifetime** [*seconds* | *infinite*]
 - **no valid-lifetime**
 - **reachable-time** *milliseconds*
 - **no reachable-time**
 - **retransmit-time** *milliseconds*
 - **no retransmit-time**
 - **router-lifetime** *seconds*
 - **router-lifetime no-default-router**
 - **no router-lifetime**
 - [no] **urpf-check**
 - **mode** {*strict* | *loose* | *strict-no-ecmp*}
 - **no mode**
 - [no] **dhcp6**
 - [no] **proxy-server**
 - **client-applications** [*dhcp*] [*ppp*]
 - **no client-applications**
 - **preferred-lifetime** *infinite*
 - **preferred-lifetime** [*days days*] [*hrs hours*] [*min minutes*] [*sec seconds*]
 - **no preferred-lifetime**
 - **rebind-timer** [*days days*] [*hrs hours*] [*min minutes*] [*sec seconds*]
 - **no rebind-timer**
 - **renew-timer** [*days days*] [*hrs hours*] [*min minutes*] [*sec seconds*]
 - **no renew-timer**
 - [no] **shutdown**
 - **valid-lifetime** *infinite*
 - **valid-lifetime** [*days days*] [*hrs hours*] [*min minutes*] [*sec seconds*]
 - **no valid-lifetime**
 - [no] **mac** *ieee-address*
 - [no] **oper-up-while-empty**
 - [no] **pppoe**
 - **description** *description-string*
 - **no description**
 - **dhcp-client**
 - [no] **ccag-use-origin-sap**
 - **pap-chap-user-db** *local-user-db-name*
 - **no pap-chap-user-db**
 - **policy** *pppoe-policy-name*
 - **no policy**
 - **sap-session-limit** *sap-session-limit*
 - **no sap-session-limit**
 - **session-limit** *session-limit*
 - **no session-limit**

- [no] **shutdown**
- **user-db** *local-user-db-name*
- **no user-db**
- **redundant-interface** *red-ip-int-name*
- **no redundant-interface**
- **shcv-policy-ipv4** *policy-name*
- **no shcv-policy-ipv4**
- **tos-marking-state** {**trusted** | **untrusted**}
- **no tos-marking-state**
- [no] **urpf-check**
 - **mode** {**strict** | **loose** | **strict-no-ecmp**}

2.5.1.9.2 Group Interface SAP Commands

- ```

config
 — service
 — ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
 — [no] subscriber-interface ip-int-name
 — group-interface ip-int-name [create]
 — no group-interface ip-int-name
 — [no] sap sap-id
 — accounting-policy acct-policy-id
 — no accounting-policy [acct-policy-id]
 — anti-spoof {ip | ip-mac | nh-mac}}
 — no anti-spoof
 — app-profile app-profile-name
 — no app-profile
 — atm
 — egress
 — traffic-desc traffic-desc-profile-id
 — no traffic-desc
 — encapsulation atm-encap-type
 — ingress
 — traffic-desc traffic-desc-profile-id
 — no traffic-desc
 — oam
 — [no] alarm-cells
 — [no] periodic-loopback
 — calling-station-id calling-station-id
 — no calling-station-id
 — [no] collect-stats
 — cpu-protection policy-id [mac-monitoring] | [eth-cfm-monitoring [aggregate] [car]] | [ip-src-monitoring]
 — no cpu-protection
 — default-host ip-address/mask next-hop next-hop-ip
 — no default-host ip-address/mask
 — description long-description-string
 — no description
 — dist-cpu-protection policy-name
 — no dist-cpu-protection
 — egress

```



- [no] **agg-rate**
  - **rate** {max | rate}
  - **no rate**
  - [no] **limit-unused-bandwidth**
  - [no] **queue-frame-based-accounting**
- **filter ip** *ip-filter-id*
- **no filter**
- **filter ipv6** *ipv6-filter-id*
- **no filter** [*ip ip-filter-id*] [*ipv6 ipv6-filter-id*]
- **policer-control-policy** *policy-name*
- **no policer-control-policy**
- [no] **qinq-mark-top-only**
- **qos** *policy-id* [**port-redirect-group** *queue-group-name* **instance** *instance-id*]
- **no qos**
- **scheduler-policy** *scheduler-policy-name*
- **no scheduler-policy**
- [no] **fwd-wholesale**
  - **pppoe** *service-id*
  - **no pppoe**
- **host ip** *ip-address* [**mac** *ieee-address*] [**subscriber** *sub-ident-string*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**ancp-string** *ancp-string*]
- **no host** {[*ip ip-address*] [*mac ieee-address*]}
- **no host all**
- **igmp-host-tracking**
  - [no] **disable-router-alert-check**
  - **expiry-time** *expiry-time*
  - **no expiry-time**
  - **import** *policy-name*
  - **no import**
  - **max-num-group** *max-num-groups*
  - **no max-num-group**
  - **max-num-sources** *max-num-sources*
  - **no max-num-sources**
  - **max-num-grp-sources** [1..32000]
  - **no max-num-grp-sources**
- [no] **shutdown**
- **ingress**
  - **filter ip** *ip-filter-id*
  - **no filter**
  - **filter ipv6** *ipv6-filter-id*
  - **no filter** [*ip ip-filter-id*] [*ipv6 ipv6-filter-id*]
  - **match-qinq-dot1p** {top | bottom}
  - **no match-qinq-dot1p**
  - **qos** *policy-id* [**shared-queuing** | **multipoint-shared**] [**fp-redirect-group** *queue-group-name* **instance** *instance-id*]
  - **no qos**
  - **scheduler-policy** *scheduler-policy-name*
  - **no scheduler-policy**
- **lag-link-map-profile** *lag-link-map-profile-id*
- **no lag-link-map-profile**
- **lag-per-link-hash** **class** {1 | 2 | 3} **weight** *weight*

- 
- **no lag-per-link-hash**
  - **multi-service-site** *customer-site-name*
  - **no multi-service-site**
  - **static-host ip** *ip/did-address* [*mac ieee-address*]  
[**create**]
  - **static-host mac** *ieee-address* [**create**]
  - **no static-host** [*ip ip-address*] **mac** *ieee-address*
  - **no static-host all** [**force**]
  - **no static-host ip** *ip-address*
    - **ancp-string** *ancp-string*
    - **no ancp-string**
    - **app-profile** *app-profile-name*
    - **no app-profile**
    - **inter-dest-id** *intermediate-destination-id*
    - **no inter-dest-id**
    - **managed-routes**
      - **route** {*ip-prefix/length* | *ip-prefix netmask*}  
[**create**]
      - **no route** {*ip-prefix/length* | *ip-prefix netmask*}
  - [**no**] **shutdown**
  - **sla-profile** *sla-profile-name*
  - **no sla-profile**
  - **sub-profile** *sub-profile-name*
  - **no sub-profile**
  - **subscriber** *sub-ident*
  - **no subscriber**
  - [**no**] **subscriber-sap-id**
  - [**no**] **shutdown**
  - [**no**] **sub-sla-mgmt**
    - **def-sla-profile** *default-sla-profile-name*
    - **no def-sla-profile**
    - **def-sub-profile** *default-subscriber-profile-name*
    - **no def-sub-profile**
    - **multi-sub-sap** *subscriber-limit*
    - **no multi-sub-sap**
    - [**no**] **shutdown**
    - **single-sub-parameters**
      - **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name*  
[**subscriber** *sub-ident-string*]
      - **no non-sub-traffic**
      - [**no**] **profiled-traffic-only**
      - **sub-ident-policy** *sub-ident-policy-name*
      - **no sub-ident-policy**
  - [**no**] **shutdown**
  - [**no**] **srrp** *srrp-id*
    - [**no**] **bfd-enable** *service-id* **interface** *interface-name* **dst-**  
**ip** *ip-address*
    - [**no**] **bfd-enable** **interface** *interface-name* **dst-ip** *ip-*  
*address* **name** *name*
    - [**no**] **bfd-enable** **interface** *interface-name* **dst-ip** *ip-*  
*address*
    - **description** *description-string*

- **no description**
- **gw-mac** *mac-address*
- **no gw-mac**
- **keep-alive-interval** *interval*
- **no keep-alive-interval**
- **message-path** *sap-id*
- **no message-path**
- **[no] policy** *vrrp-policy-id*
- **priority** *priority*
- **no priority**
- **[no] shutdown**

## Group Interface SAP ETH-CFM Commands

```
config>service>ies>sub-if>grp-if>sap
— eth-cfm
 — [no] collect-lmm-stats
 — collect-lmm-fc-stats
 — fc fc-name [fc-name]
 — no fc
 — fc-in-profile fc-name [fc-name]
 — no fc-in-profile
 — 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
 — [no] ccm-enable
 — ccm-ltm-priority priority
 — no ccm-ltm-priority
 — [no] description
 — [no] eth-test-enable
 — [no] test-pattern {all-zeros | all-ones} [crc-enable]
 — fault-propagation-enable {use-if-tlv | suspend-ccm}
 — no fault-propagation-enable
 — grace
 — eth-ed
 — max-rx-defect-window seconds
 — no max-rx-defect-window
 — priority priority
 — no priority
 — [no] rx-eth-ed
 — [no] tx-eth-ed
 — eth-vsm-grace
 — [no] rx-eth-vsm-grace
 — [no] tx-eth-vsm-grace
 — low-priority-defect {allDef | macRemErrXcon | remErrXcon | errXcon | xcon | noXcon}
 — low-priority-defect seconds
 — [no] shutdown
 — [no] squelch-ingress-levels [md-level [md-level...]]
 — tunnel-fault {accept | ignore}
```

## 2.5.2 Command Descriptions

### 2.5.2.1 Generic Commands

#### shutdown

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>        | <b>[no] shutdown</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>       | <pre> config&gt;service&gt;ies config&gt;service&gt;ies&gt;aarp-interface config&gt;service&gt;ies&gt;aarp-interface&gt;spoke-sdp config&gt;service&gt;ies&gt;if&gt;sap&gt;eth-cfm config&gt;service&gt;ies&gt;sub-if config&gt;service&gt;ies&gt;sub-if&gt;grp-if config&gt;service&gt;ies&gt;sub-if&gt;dhcp config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;dhcp config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;sap config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;sap&gt;sub-sla-mgmt config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;srrp config&gt;service&gt;ies&gt;if config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;ipv6&gt;dhcp6&gt;proxy-server config&gt;service&gt;ies&gt;if&gt;vrrp config&gt;service&gt;ies&gt;if&gt;dhcp config&gt;service&gt;ies&gt;if&gt;dhcp&gt;proxy-server config&gt;service&gt;ies&gt;if&gt;sap&gt;static-host config&gt;service&gt;ies&gt;if&gt;spoke-sdp&gt;control-channel-status config&gt;service&gt;ies&gt;redundant-interface config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;pppoe </pre> |
| <b>Description</b>   | <p>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.</p> <p>Services are created in the administratively down (<b>shutdown</b>) state. When a <b>no shutdown</b> command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.</p> <p>The <b>no</b> form of this command places the entity into an administratively enabled state.</p>                                                                                                                                                                                                                                                          |
| <b>Special Cases</b> | <p><b>IES</b> — The default administrative status of an IES service is down. While the service is down, all its associated virtual router interfaces will be operationally down. The administrative state of the service is not reflected in the administrative state of the virtual router interface.</p> <p>For example if:</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

- 1) An IES service is operational and an associated interface is shut down.
- 2) The IES service is administratively shutdown and brought back up.
- 3) The interface shutdown will remain in administrative shutdown state.

A service is regarded as operational provided that one IP Interface is operational.

Shutting down a subscriber interface will operationally shut down all child group interfaces and SAPs. Shutting down a group interface will operationally shut down all SAPs that are part of that group-interface.

**IES IP Interfaces** — When the IP interface is shutdown, it enters the administratively and operationally down states. For a SAP bound to the IP interface, no packets are transmitted out the SAP and all packets received on the SAP will be dropped while incrementing the packet discard counter.

## description

|                    |                                                                                                                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>description</b> <i>description-string</i><br><b>no description</b> [ <i>description-string</i> ]                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies<br>config>service>ies>if>dhcp<br>config>service>ies>if>sap>ip-tunnel<br>config>service>ies>sub-if>dhcp<br>config>service>ies>sub-if>grp-if>dhcp<br>config>service>ies>sub-if>grp-if>srrp<br>config>service>ies>sub-if>grp-if>pppoe                                                                                          |
| <b>Description</b> | This command creates a text description stored in the configuration file for a configuration context.<br><br>The <b>description</b> command associates a text string with a configuration context to help identify the content in the configuration file.<br><br>The <b>no</b> form of this command removes the string from the configuration. |
| <b>Default</b>     | no description                                                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <i>string</i> — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.                                                         |

## description

|                |                                                                            |
|----------------|----------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>description</b> <i>long-description-string</i><br><b>no description</b> |
| <b>Context</b> | config>service>ies>aarp-interface                                          |

```
config>service>ies>aarp-interface>spoke-sdp
config>service>ies>if>sap>eth-cfm>mep
config>service>ies>if>spoke-sdp
config>service>ies>if>spoke-sdp>eth-cfm>mep
```

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

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

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

**Default** no description

**Parameters** *long-description-string* — The description character string. Allowed values are any string up to 160 characters long 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.

## description

**Syntax** **description** *long-description-string*  
**no description** [*long-description-string*]

**Context** config>service>ies>interface  
config>service>ies>interface>sap  
config>service>ies>redundant-interface  
config>service>ies>sub-if  
config>service>ies>sub-if>grp-if  
config>service>ies>sub-if>grp-if>sap

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

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

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

**Default** no description

**Parameters** *long-description-string* — The description character string. Allowed values are any string up to 160 characters long 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.

## 2.5.2.2 IES Global Commands

### ies

|                    |                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ies</b> <i>service-id</i> [ <b>customer</b> <i>customer-id</i> ] [ <b>create</b> ] [ <b>vpn</b> <i>vpn-id</i> ] [ <b>name</b> <i>name</i> ]<br><b>no ies</b> <i>service-id</i> |
| <b>Context</b>     | config>service                                                                                                                                                                    |
| <b>Description</b> | This command creates or edits an IES service instance.                                                                                                                            |

The **ies** command is used to create or maintain an Internet Enhanced Service (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.

IES services allow the creation of customer facing IP interfaces in the same routing instance used for service network core routing connectivity. IES 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.

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 set aside for service IP provisioning, becoming administered by a separate but subordinate address authority. This feature is defined using the **config router service-prefix** command.

IP interfaces defined within the context of an IES service ID must have a SAP created as the access point to the subscriber network. This allows a combination of bridging and IP routing for redundancy purposes.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. Once a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

Once a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

Multiple IES services are created to separate customer owned IP interfaces. More than one IES service may be created for a single customer ID. More than one IP interface may be created within a single IES service ID. All IP interfaces created within an IES service ID belongs to the same customer.

By default, no IES service instances exist until they are explicitly created.

The **no** form of this command deletes the IES service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces defined within the service ID have been shutdown and deleted.

---

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <p><i>service-id</i> — 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 <i>service-id</i> must be the same number used for every SR OS on which this service is defined.</p> <p><b>Values</b>     <i>service-id</i>: 1 to 2147483648<br/>                  <i>svc-name</i>: 64 characters maximum</p> <p><b>customer</b> <i>customer-id</i> — 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.</p> <p><b>Values</b>     1 to 2147483647</p> <p><b>vpn</b> <i>vpn-id</i> — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.</p> <p><b>Values</b>     1 to 2147483647</p> <p><b>Default</b>     null (0)</p> <p><b>name</b> <i>name</i> — Configures an optional service name identifier, up to 64 characters, to a given service. This service name can then be used in configuration references, display, and show commands throughout the system. A defined service name can help the service provider or administrator to identify and manage services within the SR OS platforms.</p> <p>To create a service, you must assign a service ID; however, after it is created, either the service ID or the service name can be used to identify and reference a service. If a name is not specified at creation time, then SR OS assigns a string version of the <i>service-id</i> as the name.</p> <p><b>Values</b>     <i>name</i>: 64 characters maximum</p> |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## igmp-host-tracking

|                    |                                                                             |
|--------------------|-----------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>igmp-host-tracking</b>                                                   |
| <b>Context</b>     | config>service>ies<br>config>service>ies>sub-if>grp-if>sap                  |
| <b>Description</b> | This command enters the context to configure IGMP host tracking parameters. |

## disable-router-alert-check

|                    |                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] disable-router-alert-check</b>                                                                                                    |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap>igmp-host-tracking                                                                                   |
| <b>Description</b> | <p>This command disables the IGMP router alert check option.</p> <p>The <b>no</b> form of the command enables the router alert check.</p> |



## expiry-time

|                    |                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>expiry-time</b> <i>expiry-time</i><br><b>no expiry-time</b>                                                                                                          |
| <b>Context</b>     | config>service>ies>igmp-host-tracking<br>config>service>ies>sub-if>grp-if>sap>igmp-host-tracking                                                                        |
| <b>Description</b> | This command configures the time that the system continues to track inactive hosts.<br><br>The <b>no</b> form of the command removes the values from the configuration. |
| <b>Default</b>     | no expiry-time                                                                                                                                                          |
| <b>Parameters</b>  | <i>expiry-time</i> — Specifies the time, in seconds, that this system continues to track an inactive host.<br><br><b>Values</b> 1 to 65535                              |

## max-num-group

|                    |                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>max-num-groups</b> <i>max-num-groups</i><br><b>no max-num-groups</b>                                                                            |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap>igmp-host-tracking                                                                                            |
| <b>Description</b> | This command configures the maximum number of multicast groups allowed to be tracked.<br><br>The <b>no</b> form of the command disables the check. |
| <b>Default</b>     | no max-num-groups                                                                                                                                  |
| <b>Parameters</b>  | <i>max-num-groups</i> — Specifies the maximum number of multicast groups allowed to be tracked.<br><br><b>Values</b> 1 to 196607                   |

## max-num-sources

|                    |                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>max-num-sources</b> <i>max-num-sources</i><br><b>no max-num-sources</b>                                                                                                   |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap>igmp-host-tracking                                                                                                                      |
| <b>Description</b> | This command configures the maximum number of multicast sources allowed to be tracked per group.<br><br>The no form of the command removes the value from the configuration. |

---

|                   |                                                                                                             |
|-------------------|-------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>max-num-sources</i> — Specifies the maximum number of multicast sources allowed to be tracked per group. |
| <b>Values</b>     | 1 to 1000                                                                                                   |

## max-num-grp-sources

|                    |                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>max-num-grp-sources</b> <i>max-num-sources</i><br><b>no max-num-grp-sources</b>                                                              |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap>igmp-host-tracking                                                                                         |
| <b>Description</b> | This command configures the max number of multicast (S,G)s allowed to be tracked.<br><br>The <b>no</b> form of this command disables the check. |
| <b>Default</b>     | no max-num-grp-sources                                                                                                                          |
| <b>Parameters</b>  | <i>max-num-sources</i> — Specifies the maximum number of multicast sources allowed to be tracked per group.                                     |
| <b>Values</b>      | 1 to 32000                                                                                                                                      |

## import

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>import</b> <i>policy-name</i><br><b>no import</b>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap>igmp-host-tracking                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | This command specifies the import routing policy to be used for IGMP packets to be used on this SAP. Only a single policy can be imported on a single SAP at any time.<br><br>The <b>no</b> form of the command removes the policy association from the SAP.                                                                                                                                                          |
| <b>Default</b>     | no import — No import policy is specified.                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <i>policy-name</i> — The import policy name. Values can be string up to 32 characters long of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. These policies are configured in the <b>config&gt;router&gt; policy-options</b> context. The router policy must be defined before it can be imported. |

### 2.5.2.3 AARP Interface Commands

#### aarp-interface

|                    |                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>aarp-interface</b> <i>aarp-interface-name</i> [ <b>create</b> ]<br><b>no aarp-interface</b> <i>aarp-interface-name</i>                                                                                                                                                                                                                                |
| <b>Context</b>     | config>service>ies                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | This command creates an AARP interface for connecting a service to a peer node AARP service. This instance is paired with the same AARP interface in a peer node as part of a configuration to provide flow and packet asymmetry removal for traffic for a multi-homed SAP or spoke SDP.<br><br>The <b>no</b> form of the command deletes the interface. |
| <b>Default</b>     | no aarp-interface                                                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <i>aarp-interface-name</i> — A string of up to 32 characters identifying the interface.<br><b>create</b> — Keyword used to create the AARP interface.                                                                                                                                                                                                    |

#### ip-mtu

|                    |                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ip-mtu</b> <i>octets</i><br><b>no ip-mtu</b>                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>aarp-interface                                                                                                                                                                                                                      |
| <b>Description</b> | This command configures the IP maximum transmit unit (packet) for this interface.<br><br>The <b>no</b> form of the command returns the default value. By default (for Ethernet network interface) if no ip-mtu is configured it is (1568 - 14) = 1554. |
| <b>Default</b>     | no ip-mtu                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <i>octets</i> — Specifies the maximum number of octets that can be transmitted.<br><br><b>Values</b> 512 to 9000                                                                                                                                       |

#### spoke-sdp

|                |                                                                                                   |
|----------------|---------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>spoke-sdp</b> <i>sdp-id:vc-id</i> [ <b>create</b> ]<br><b>no spoke-sdp</b> <i>sdp-id:vc-id</i> |
| <b>Context</b> | config>service>ies>aarp-interface                                                                 |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command binds a service to an existing SDP. A spoke SDP is treated like the equivalent of a traditional bridge port where flooded traffic received on the spoke SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.</p> <p>The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.</p> <p>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.</p> <p>The <b>no</b> form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.</p> |
| <b>Default</b>     | no spoke-sdp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><i>sdp-id</i> — Specifies the SDP identifier.</p> <p><b>Values</b> 1 to 17407</p> <p><i>vc-id</i> — The virtual circuit identifier. The VC-ID is not used with L2TPv3 SDPs, however it must be configured.</p> <p><b>Values</b> 1 to 4294967295</p> <p><b>create</b> — Keyword used to create the spoke SDP.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## aarp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>aarp <i>aarp-id</i> type {subscriber-side-shunt   network-side-shunt}</b><br><b>no aarp</b>                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>service>ies>aarp-interface>spoke-sdp                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command associates an AARP instance to an AARP interface spoke SDP. This instance is paired with the same <i>aarp-id</i> in a peer node as part of a configuration to provide flow and packet asymmetry removal for traffic for a multi-homed SAP or spoke SDP. The <b>type</b> parameter specifies the role of this service point in the AARP instance.</p> <p>The <b>no</b> form of the command removes the association.</p> |
| <b>Default</b>     | no aarp                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><i>aarp-id</i> — An integer that identifies an AARP instance.</p> <p><b>Values</b> 1 to 65535</p> <p><b>subscriber-side-shunt</b> — Specifies that the AARP type is an inter-chassis shunt service for subscriber-side traffic.</p>                                                                                                                                                                                                 |

**network-side-shunt** — Specifies that the AARP type is an inter-chassis shunt service for network-side traffic.

## egress

|                    |                                                         |
|--------------------|---------------------------------------------------------|
| <b>Syntax</b>      | <b>egress</b>                                           |
| <b>Context</b>     | config>service>ies>aarp-interface>spoke-sdp             |
| <b>Description</b> | This command enters the egress context for a spoke SDP. |

## filter

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>filter ip</b> <i>ip-filter-id</i><br><b>no filter</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>service>ies>aarp-interface>spoke-sdp>egress<br>config>service>ies>aarp-interface>spoke-sdp>ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command associates an IP filter policy with an ingress or egress IP interface. Filter policies control the forwarding and dropping of packets based on IP matching criteria.</p> <p>The <i>filter-id</i> must already be defined before the filter command is executed. If the filter policy does not exist, the operation fails and an error message returned.</p> <p>IP filters apply only to RFC 2427-routed IP packets. Frames that do not contain IP packets will not be subject to the filter and will always be passed, even if the filter's default action is to drop.</p> <p>The <b>no</b> form of this command removes any configured filter ID association with the IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local.</p> |
| <b>Parameters</b>  | <p><i>ip-filter-id</i> — Specifies the filter policy. The filter ID must already exist within the created IP filters.</p> <p><b>Values</b> 1 to 65535 or a string up to 64 characters</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## vc-label

|                    |                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>vc-label</b> <i>vc-label</i><br><b>no vc-label</b> [ <i>vc-label</i> ]                                 |
| <b>Context</b>     | config>service>ies>aarp-interface>spoke-sdp>egress<br>config>service>ies>aarp-interface>spoke-sdp>ingress |
| <b>Description</b> | This command configures the egress and ingress VC label.                                                  |

The **no** version of this command removes the VC label.

|                   |                                                                           |
|-------------------|---------------------------------------------------------------------------|
| <b>Parameters</b> | <i>vc-label</i> — A VC egress value that indicates a specific connection. |
| <b>Values</b>     | egress: 16 to 1048575<br>ingress: 32 to 18431                             |

## ingress

|                    |                                                          |
|--------------------|----------------------------------------------------------|
| <b>Syntax</b>      | <b>ingress</b>                                           |
| <b>Context</b>     | config>service>ies>aarp-interface>spoke-sdp              |
| <b>Description</b> | This command enters the ingress context for a spoke SDP. |

### 2.5.2.4 IES Interface Commands

## interface

|                    |                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>interface</b> <i>ip-int-name</i> [ <b>create</b> ]<br><b>interface</b> <i>ip-int-name</i> [ <b>create</b> ] <b>tunnel</b><br><b>no interface</b> <i>ip-int-name</i>                                          |
| <b>Context</b>     | config>service>ies                                                                                                                                                                                              |
| <b>Description</b> | This command creates a logical IP routing interface for an Internet Enhanced Service (IES). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface. |

The **interface** command, under the context of services, 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 table. The typical use for IP interfaces created in this manner is for subscriber Internet access. An IP address cannot be assigned to an IES interface. Multiple SAPs can be assigned to a single group interface.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config router interface** and **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.

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 IES IP interfaces. It also controls the prefixes that may be learned or statically defined with the IES 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 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 shutdown before the SAP on that interface may be removed. IES services do not have the **shutdown** command in the SAP CLI context. IES 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 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 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.
  - tunnel** — Specifies that this is an IPsec interface used for IPsec tunneling.
  - create** — Keyword used to create the interface. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

## address

- Syntax**
- address** {*ip-address/mask* | *ip-address netmask*} [**broadcast** {**all-ones** | **host-ones**}] [**track-srrp** *srrp-instance*]
- no address** [*ip-address/mask* | *ip-address netmask*]
- Context**
- config>service>ies>if  
config>service>ies>subscriber-interface

**Description** This command assigns an IP address, IP subnet, and broadcast address format 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 router.

For the 7750 SR only, in the IES subscriber interface context, this command is used to assign one or more host IP addresses and subnets. This differs from a normal IES interfaces where the **secondary** command creates an additional subnet after the primary address is assigned. A user can then add or remove addresses without having to keep a primary address.

The local subnet that the **address** command defines must be part of the services address space within the routing context using the **config router service-prefix** command. The default is to disallow the complete address space to services. Once a portion of the address space is allocated as a service prefix, that portion can be made unavailable for IP interfaces defined within the **config router interface** CLI context for network core connectivity with the **exclude** option in the **config router service-prefix** command.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. The show commands display CIDR notation and is stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

Use the **no** form of this command to remove the IP address assignment from the IP interface.

The **no** form of this command will cause ptp-hw-assist to be disabled.

**Table 7 Address Field Descriptions**

| 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 and the protocol interfaces and the MPLS LSPs associated with that IP interface will be reinitialized.



- 
- Parameters**
- ip-address* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).
  - /* — The forward slash is a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the “/” and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted decimal mask must follow the prefix.
  - mask-length* — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 0 – 30. A mask length of 32 is reserved for system IP addresses.
  - mask* — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical ‘AND’ function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. A mask of 255.255.255.255 is reserved for system IP addresses.
  - netmask* — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.
  - broadcast** — The optional **broadcast** parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**. The broadcast format on an IP interface can be specified when the IP address is assigned or changed.  
  
This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) will be received by the IP interface. (*Default: host-ones*)
  - all-ones** — The **all-ones** keyword following the **broadcast** parameter specifies the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.
  - host-ones** — The **host-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default broadcast address used by an IP interface.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

**track-srrp** — Specifies the SRRP instance ID that this interface route needs to track.

## address

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>address</b> <i>ip-address/mask</i> [ <i>netmask</i> ] [ <b>gw-ip-address</b> <i>ip-address</i> ] [ <b>populate-host-routes</b> ] [ <b>track-srrp</b> <i>srrp-instance</i> ] [ <b>holdup-time</b> <i>milli-seconds</i> ]<br><b>no address</b> <i>ip-address/mask</i> [ <i>netmask</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>service>ies>subscriber-interface<br>config>service>vprn>subscriber-interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command configures the local subscriber subnets available on a subscriber IP interface. The configured ip-address and mask define the address space associated with the subscriber subnet. Each subnet supports a locally owned IP host address within the subnet that is not expected to appear on other routers that may be servicing the same subscriber subnet. For redundancy purposes, the keyword <b>gw-address</b> defines a separate IP address within the subnet for Subscriber Routed Redundancy Protocol (SRRP) routing. This IP address must be the same on the local and remote routers participating in a common SRRP instance.</p> <p>In SRRP, a single SRRP instance is tied to a group IP interface. The group IP interface is contained directly within a subscriber IP interface context and thus directly associated with the subscriber subnets on the subscriber IP interface. The SRRP instance is also indirectly associated with any subscriber subnets tied to the subscriber interface through wholesale/retail VPRN configurations. With the directly-associated and the indirectly-associated subscriber interface subnets, a single SRRP instance can manage hundreds of SRRP gateway IP addresses. This automatic subnet association to the SRRP instance is different from VRRP where the redundant IP address is defined within the VRRP context.</p> <p>Defining an SRRP gateway IP address on a subscriber subnet is not optional when the subnet is associated with a group IP interface with SRRP enabled. Enabling SRRP (<b>no shutdown</b>) fails if one or more subscriber subnets do not have an SRRP gateway IP address defined. Creating a new subscriber subnet without an SRRP gateway IP address defined fails when the subscriber subnet is associated with a group IP interface with an active SRRP instance. Once SRRP is enabled on a group interface, the SRRP instance will manage the ARP response and routing behavior for all subscriber hosts reachable through the group IP interface.</p> <p>The no form of the command removes the address from a subscriber subnet. The <b>address</b> command for the specific subscriber subnet must be executed without the <b>gw-address</b> parameter. To succeed, all SRRP instances associated with the subscriber subnet must be removed or shutdown.</p> |
| <b>Parameters</b>  | <i>ip-address/mask</i>   <i>ip-address netmask</i> — Specifies the address space associated with the subscriber subnet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**gw-ip-address** *ip-address* — Specifies a separate IP address within the subnet for SRRP routing purposes. This parameter must be followed by a valid IP interface that exists within the subscriber subnet created by the address command. The defined gateway IP address cannot currently exist as a subscriber host (static or dynamic). If the defined ip-address already exists as a subscriber host address, the address command fails. The specified ip-address must be unique within the system.

The gw-address parameter may be specified at anytime. If the subscriber subnet was created previously, executing the address command with a gw-address parameter will simply add the SRRP gateway IP address to the existing subnet.

If the address command is executed without the gw-address parameter when the subscriber subnet is associated with an active SRRP instance, the address fails. If the SRRP instance is inactive or removed, executing the address command without the gw-address parameter will remove the SRRP gateway IP address from the specified subscriber subnet.

If the address command is executed with a new gw-address, all SRRP instances currently associated with the specified subscriber subnet will be updated with the new SRRP gateway IP address.

**populate-host-routes** — Indicates that all subscriber-hosts created on the interface with the ip-address falling in this subnet will have their route populated in FIB. This flag will not be set per default.

**track-srrp** *srrp-instance* — Specifies the SRRP instance ID that this interface route needs to track.

**Values** 1 to 4294967295

**holdup-time** *milli-seconds* — Specifies the holdup time, in milliseconds.

**Values** 100 to 5000

## allow-unmatching-subnets

|                    |                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] allow-unmatching-subnets</b>                                                                                                            |
| <b>Context</b>     | config>service>ies>subscriber-interface                                                                                                         |
| <b>Description</b> | This command specifies whether subscriber hosts with a subnet that does not match any of the subnets configured on this interface, are allowed. |

## allow-directed-broadcasts

|                    |                                                                                     |
|--------------------|-------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] allow-directed-broadcasts</b>                                               |
| <b>Context</b>     | config>service>ies>if                                                               |
| <b>Description</b> | This command enables the forwarding of directed broadcasts out of the IP interface. |

A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address on another IP interface. The **allow-directed-broadcasts** command on an IP interface enables or disables the transmission of packets destined to the subnet broadcast address of the egress IP interface.

When enabled, a frame destined to the local subnet on this IP interface will be sent as a subnet broadcast out this interface. Care should be exercised when allowing directed broadcasts as it is a well-known mechanism used for denial-of-service attacks.

When disabled, directed broadcast packets discarded at this egress IP interface will be counted in the normal discard counters for the egress SAP.

By default, directed broadcasts are not allowed and will be discarded at this egress IP interface.

The **no** form of this command disables the forwarding of directed broadcasts out of the IP interface.

**Default** no allow-directed-broadcasts — Directed broadcasts are dropped.

## arp-learn-unsolicited

**Syntax** [no] arp-learn-unsolicited

**Context** config>service>ies>if

**Description** This command allows the ARP application to learn new entries based on any received ARP message (GARP, ARP-Request, or ARP-Reply, such as any frame with ethertype 0x0806).

The **no** form of this command disables the above behavior and causes ARP entries to only be learned when needed, that is, when the router receives an ARP-reply after an ARP-request triggered by received traffic.

## arp-limit

**Syntax** arp-limit *limit* [log-only] [threshold *percent*]  
no arp-limit

**Context** config>service>ies>interface

**Description** This command configures the maximum amount of dynamic IPv4 ARP entries that can be learned on an IP interface.

When the number of dynamic ARP entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations will be dropped. Entries that have already been learned will be refreshed.

The **no** form of the command removes the **arp-limit**.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no arp-limit                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b> | <p><b>log-only</b> — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit will be learned.</p> <p><i>percent</i> — The threshold value (as a percentage) that triggers a warning message to be sent.</p> <p><b>Values</b> 0 to 100</p> <p><i>limit</i> — The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic ARP learning is disabled and no dynamic ARP entries are learned.</p> <p><b>Values</b> 0 to 524288</p> |

## arp-populate

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] arp-populate</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>service>ies>if<br>config>service>ies>sub-if>grp-if                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command, when enabled, disables dynamic learning of ARP entries. Instead, the ARP table is populated with dynamic entries from the DHCP Lease State Table (enabled with <b>lease-populate</b>), and optionally with static entries entered with the <b>host</b> command.</p> <p>Enabling the <b>arp-populate</b> command will remove any dynamic ARP entries learned on this interface from the ARP cache.</p> <p>The <b>arp-populate</b> command fails if an existing static ARP entry exists for this interface. The <b>arp-populate</b> command fails if an existing static subscriber host on the SAP does not have both MAC and IP addresses specified.</p> <p>Once <b>arp-populate</b> is enabled, creating a static subscriber host on the SAP without both an IP address and MAC address fails.</p> <p>When <b>arp-populate</b> is enabled, the system will not send out ARP requests for hosts that are not in the ARP cache. Only statically configured and DHCP learned hosts are reachable through an IP interface with <b>arp-populate</b> enabled. The <b>arp-populate</b> command can only be enabled on IES and VPRN interfaces supporting Ethernet encapsulation.</p> <p>Use the <b>no</b> form of the command to disable ARP cache population functions for static and dynamic hosts on the interface. All static and dynamic host information for this interface will be removed from the system's ARP cache.</p> |
| <b>Default</b>     | not enabled                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

---

## arp-populate-host-route

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] arp-populate-host-route</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>service>ies>if                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command enables the addition or deletion of host routes in the route table derived from ARP entries in the ARP cache. To enable this command, the interface must be shut down. The command triggers the population of host routes in the route table out of their corresponding static, dynamic, or EVPN types in the ARP table. ARP entries installed by subscriber management, local interfaces, and others, do not create host routes.</p> <p>The <b>no</b> form of this command disables the creation of host routes from the ARP cache.</p> |

## arp-proactive-refresh

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] arp-proactive-refresh</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>service>ies>if                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command enables the router to always send out a single refresh message with no entries 30 seconds prior to the timeout of the entry.</p> <p>The <b>no</b> form of this command sets the default behavior, in which an entry is marked as stale 30 seconds prior to age-out, and the router only sends an ARP request to refresh the entry if the IOM receives traffic that uses it. If so, the IOM asks the ARP application to send a refresh message. With <b>arp-proactive-refresh</b> enabled, the ARP module sends a refresh message regardless of whether the IOM receives traffic.</p> |

## arp-retry-timer

|                    |                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>arp-retry-timer <i>timer-multiple</i></b><br><b>no arp-retry-timer</b>                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>if                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command allows the arp retry timer to be configured to a specific value.</p> <p>The timer value is entered as a multiple of 100 ms. So a timer value of 1, means the ARP timer will be set to 100 ms.</p> <p>The <b>no</b> form of this command removes the command from the active configuration and returns the ARP retry timer to its default value of 5 seconds.</p> |
| <b>Default</b>     | arp-retry-timer 50                                                                                                                                                                                                                                                                                                                                                               |

---

|                   |                                                                                                          |
|-------------------|----------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>timer-multiple</i> — Specifies the multiple of 100 ms that the ARP retry timer will be configured as. |
| <b>Values</b>     | 1 to 300 (equally a timer range of 100 ms to 30,000 ms)                                                  |

## arp-route-tag

|                    |                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>arp-route-tag</b> <i>tag</i><br><b>arp-route-tag</b>                                                                                                                                                                                                              |
| <b>Context</b>     | config>service>ies>if                                                                                                                                                                                                                                                |
| <b>Description</b> | This command adds a route tag to the ARP-ND host routes generated from the ARP entries in the interface which can be used to match ARP-ND routes in BGP export policies.<br><br>The <b>no</b> form of this command removes the route tag for the ARP-ND host routes. |
| <b>Parameters</b>  | <i>tag</i> — Specifies the route tag value.                                                                                                                                                                                                                          |
| <b>Values</b>      | 1 to 255                                                                                                                                                                                                                                                             |

## arp-timeout

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>arp-timeout</b> <i>seconds</i><br><b>no arp-timeout</b> [ <i>seconds</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>service>ies>if<br>config>service>ies>sub-if>grp-if                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | 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 <b>arp-timeout</b> is set to a value of zero seconds, ARP aging is disabled.<br><br>When the <b>arp-populate</b> and <b>lease-populate</b> commands are enabled on an IES interface, the ARP table entries will no longer be dynamically learned, but instead by snooping DHCP ACK message from a DHCP server. In this case the configured <b>arp-timeout</b> value has no effect.<br><br>The <b>no</b> form of this command restores <b>arp-timeout</b> to the default value. |
| <b>Default</b>     | arp-timeout 14400                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <i>seconds</i> — 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Values</b>      | 0 to 65535                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

---

bfd

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |               |                                                                          |                |     |               |                                                                          |                |     |               |         |                |   |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|--------------------------------------------------------------------------|----------------|-----|---------------|--------------------------------------------------------------------------|----------------|-----|---------------|---------|----------------|---|
| <b>Syntax</b>      | <b>bfd</b> <i>transmit-interval</i> [ <b>receive</b> <i>receive-interval</i> ] [ <b>multiplier</b> <i>multiplier</i> ] [ <b>echo-receive</b> <i>echo-interval</i> ] [ <b>type</b> <b>cpm-np</b> ]<br><b>no bfd</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |               |                                                                          |                |     |               |                                                                          |                |     |               |         |                |   |
| <b>Context</b>     | config>service>ies>if<br>config>service>ies>if>ipv6                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |               |                                                                          |                |     |               |                                                                          |                |     |               |         |                |   |
| <b>Description</b> | <p>This command specifies the BFD parameters for the associated IP interface. If no parameters are defined the default value are used.</p> <p>The multiplier specifies the number of consecutive BFD messages that must be missed from the peer before the BFD session state is changed to down and the upper level protocols (OSPF, IS-IS, BGP or PIM) is notified of the fault.</p> <p>The <b>no</b> form of the command removes BFD from the interface.</p> <div style="margin-top: 20px;">  <p><b>Note:</b> On the 7750 SR, the <i>transmit-interval</i>, <b>receive</b> <i>receive-interval</i>, and <b>echo-receive</b> <i>echo-interval</i> values can only be modified to a value less than 100 when:</p> <ol style="list-style-type: none"> <li>1. The <b>type cpm-np</b> option is explicitly configured.</li> <li>2. The service is shut down (<b>shutdown</b>)</li> <li>3. The interval is specified 10 to 100000.</li> <li>4. The service is re-enabled (<b>no shutdown</b>)</li> </ol> <p>To remove the <b>type cpm-np</b> option, re-issue the <b>bfd</b> command without specifying the <b>type</b> parameter.</p> </div> |               |                                                                          |                |     |               |                                                                          |                |     |               |         |                |   |
| <b>Default</b>     | no bfd                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |               |                                                                          |                |     |               |                                                                          |                |     |               |         |                |   |
| <b>Parameters</b>  | <p><i>transmit-interval</i> — Sets the transmit interval for the BFD session.</p> <table> <tr> <td><b>Values</b></td><td>100 to 100000<br/>10 to 100000 (for the 7750 SR only; see the Note above)</td></tr> <tr> <td><b>Default</b></td><td>100</td></tr> </table> <p><b>receive</b> <i>receive-interval</i> — Sets the receive interval for the BFD session.</p> <table> <tr> <td><b>Values</b></td><td>100 to 100000<br/>10 to 100000 (for the 7750 SR only; see the Note above)</td></tr> <tr> <td><b>Default</b></td><td>100</td></tr> </table> <p><b>multiplier</b> <i>multiplier</i> — Sets the multiplier for the BFD session.</p> <table> <tr> <td><b>Values</b></td><td>3 to 20</td></tr> <tr> <td><b>Default</b></td><td>3</td></tr> </table>                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>Values</b> | 100 to 100000<br>10 to 100000 (for the 7750 SR only; see the Note above) | <b>Default</b> | 100 | <b>Values</b> | 100 to 100000<br>10 to 100000 (for the 7750 SR only; see the Note above) | <b>Default</b> | 100 | <b>Values</b> | 3 to 20 | <b>Default</b> | 3 |
| <b>Values</b>      | 100 to 100000<br>10 to 100000 (for the 7750 SR only; see the Note above)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |               |                                                                          |                |     |               |                                                                          |                |     |               |         |                |   |
| <b>Default</b>     | 100                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |               |                                                                          |                |     |               |                                                                          |                |     |               |         |                |   |
| <b>Values</b>      | 100 to 100000<br>10 to 100000 (for the 7750 SR only; see the Note above)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |               |                                                                          |                |     |               |                                                                          |                |     |               |         |                |   |
| <b>Default</b>     | 100                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |               |                                                                          |                |     |               |                                                                          |                |     |               |         |                |   |
| <b>Values</b>      | 3 to 20                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |               |                                                                          |                |     |               |                                                                          |                |     |               |         |                |   |
| <b>Default</b>     | 3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |               |                                                                          |                |     |               |                                                                          |                |     |               |         |                |   |



**echo-receive** *echo-interval* — Sets the minimum echo receive interval, in milliseconds, for the BFD session.

**Values** 100 to 100000  
10 to 100000 (for the 7750 SR only; see the Note above)

**Default** 100

**type cpm-np** — For the 7750 SR only, specifies that BFD sessions associated with this interface will be created on the CPM network processor to allow for fast timers down to 10 ms granularity.

## cflowd-parameters

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>cflowd-parameters</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>service>ies>if<br>config>service>ies>sub-if>grp-if                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command creates the configuration context to configure cflowd parameters for the associated IP interfaces.</p> <p><b>cflowd</b> is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement.</p> <p>At a minimum, the <b>sampling</b> command must be configured within this context in order to enable cflowd sampling, otherwise traffic sampling will not occur.</p> |
| <b>Default</b>     | no cflowd-parameters                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## sampling

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sampling {unicast   multicast} type {acl   interface} [direction {ingress-only   egress-only   both}]</b><br><b>no sampling {unicast   multicast}</b>                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>service>ies>if>cflowd-parameters<br>config>service>ies>sub-if>grp-if>cflowd-parameters                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command enables and configures the cflowd sampling behavior to collect traffic flow samples through a router for analysis.</p> <p>This command can be used to configure the sampling parameters for unicast and multicast traffic separately. If sampling is not configured for either unicast or multicast traffic, then that type of traffic will not be sampled.</p> <p>Egress sampled flows are only sent to v9 or v10(IPFIX) collectors.</p> |

The **no** form of the command disables the associated type of traffic sampling on the associated interface.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no sampling                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b> | <p><b>unicast</b> — Specifies that the <b>sampling</b> command will control the sampling of unicast traffic on the associated interface/SAP.</p> <p><b>multicast</b> — Specifies that the <b>sampling</b> command will control the sampling of multicast traffic on the associated interface/SAP.</p> <p><b>type</b> — Specifies the cflowd sampling type on the specified virtual router interfaces.</p> <p><b>Values</b></p> <p>    <b>acl</b> — Specifies that the sampled traffic is controlled via an IP traffic filter entry with the action “filter-sample” configured.</p> <p>    <b>interface</b> — Specifies that all traffic entering or exiting the interface is subject to sampling.</p> <p><b>direction</b> — Specifies the direction in which to collect traffic flow samples.</p> <p><b>Values</b></p> <p>    <b>ingress-only</b> — Enables ingress sampling only on the associated interface.</p> <p>    <b>egress-only</b> — Enables egress sampling only on the associated interface.</p> <p>    <b>both</b> — Enables both ingress and egress cflowd sampling.</p> |

## cpu-protection

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>cpu-protection</b> <i>policy-id</i><br><b>no cpu-protection</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>service>ies>if                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command assigns an existing CPU protection policy to the associated service interface. For these interface types, the per-source rate limit is not applicable. The CPU protection policies are configured in the <b>config&gt;sys&gt;security&gt;cpu-protection&gt;policy</b> <i>cpu-protection-policy-id</i> context.</p> <p>If no cpu-protection policy is assigned to a service interface, then the default policy is used to limit the overall-rate. The default policy is policy number 254 for access interfaces and 255 for network interfaces.</p> <p>The no form of the command removes the association of the CPU protection policy from the associated interface and reverts to the default policy values.</p> <p>cpu-protection 254 (for access interfaces)</p> <p>cpu-protection 255 (for network interfaces)</p> <p>none (for video-interfaces, shown as no cpu-protection in CLI)</p> |

The configuration of **no cpu-protection** returns the interface/SAP to the default policies as shown above.

**Parameters** *policy-id* — Specifies an existing CPU protection policy.

**Values** 1 to 255

## cpu-protection

**Syntax** **cpu-protection** *policy-id* [**mac-monitoring**] | [**eth-cfm-monitoring** [**aggregate**][**car**]] | [**ip-src-monitoring**]  
**no cpu-protection**

**Context** config>service>ies>sub-if>grp-if>sap

**Description** This command assigns an existing CPU protection policy to the associated group interface. The CPU protection policies are configured in the **config>sys>security>cpu-protection>policy** *cpu-protection-policy-id* context.

If no CPU-Protection policy is assigned to a group interface SAP, then the default policy is used to limit the overall-rate. The default policy is policy number 254 for access interfaces and 255 for network interfaces.

The no form of the command removes the association of the CPU protection policy from the associated interface and reverts to the default policy values.

**Default** cpu-protection 254 (for access interfaces)

cpu-protection 255 (for network interfaces)

The configuration of no cpu-protection returns the interface/SAP to the default policies as shown above.

**Parameters** *policy-id* — Specifies an existing CPU protection policy.

**Values** 1 to 255

**mac-monitoring** — Enables per SAP + source MAC address rate limiting using the per-source-rate from the associated cpu-protection policy.

**eth-cfm-monitoring** — Enables Ethernet Connectivity Fault Management monitoring.

**aggregate** — Applies the rate limit to the sum of the per peer packet rates.

**car** — (Committed Access Rate) causes Eth-CFM packets to be ignored when enforcing the overall-rate.

**ip-src-monitoring** — Enables per SAP + IP source address rate limiting for DHCP packets using the per-source-rate from the associated cpu-protection policy. The ip-src-monitoring is useful in subscriber management architectures that have routers between the subscriber and the BNG (router). In Layer 3 aggregation scenarios all packets from all subscribers behind the same aggregation router will arrive with the same source MAC address and as such the mac-monitoring functionality can not differentiate traffic from different subscribers.

## ipcp

|                    |                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipcp</b>                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>service>ies>if                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | This command creates allows access to the IPCP context within the interface configuration. Within this context, IPCP extensions can be configured to define such things as the remote IP address and DNS IP address to be signaled via IPCP on the associated PPP interface. This command is only applicable if the associated SAP/port is a PPP/MLPPP interface. |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                                                                                                                               |

## dns

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dns ip-address</b> [ <b>secondary ip-address</b> ]<br><b>dns secondary ip-address</b><br><b>no dns</b> [ <i>ip-address</i> ] [ <b>secondary ip-address</b> ]                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies>if>ipcp                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | This command defines the dns address(es) to be assigned to the far-end of the associated PPP/MLPPP link through IPCP extensions. This command is only applicable if the associated SAP/port is a PPP/MLPPP interface with an IPCP encapsulation.<br><br>The <b>no</b> form of the command deletes either the specified primary DNS address, secondary DNS address or both addresses from the IPCP extension peer-ip-address configuration. |
| <b>Default</b>     | no dns                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies a unicast IPv4 address for the primary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions.<br><br><b>secondary ip-address</b> — Specifies a unicast IPv4 address for the secondary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions.                                                                                   |

---

## peer-ip-address

|                    |                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>peer-ip-address</b> <i>ip-address</i><br><b>no peer-ip-address</b>                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>service>ies>if>ipcp                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command defines the remote IP address to be assigned to the far-end of the associated PPP/MLPPP link via IPCP extensions. This command is only applicable if the associated SAP/port is a PPP/MLPPP interface with an IPCP encapsulation.</p> <p>The <b>no</b> form of the command deletes the IPCP extension peer-ip-address configuration.</p> |
| <b>Default</b>     | no peer-ip-address (0.0.0.0)                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies a unicast IPv4 address to be signaled to the far-end of the associated PPP/MLPPP link by IPCP extensions.                                                                                                                                                                                                                  |

## ipv6

|                    |                                                                        |
|--------------------|------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ipv6</b>                                                       |
| <b>Context</b>     | config>service>ies>sub-if>grp-if                                       |
| <b>Description</b> | This command enables IPv6 forwarding on the specified group-interface. |

## router-advertisements

|                    |                                                                                 |
|--------------------|---------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] router-advertisements</b>                                               |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6                                           |
| <b>Description</b> | This command enables router advertisement transmission on this group interface. |
| <b>Default</b>     | router-advertisements                                                           |

## current-hop-limit

|                    |                                                                                    |
|--------------------|------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>current-hop-limit</b> <i>hop-count</i><br><b>no current-hop-limit</b>           |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>rtr-adv                                      |
| <b>Description</b> | This command specifies the hop-limit advertised to hosts in router advertisements. |
| <b>Default</b>     | current-hop-limit 64                                                               |

---

|                   |                                                                                                   |
|-------------------|---------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>hop-count</i> — Specifies the current hop limit (decimal) inserted into router advertisements. |
| <b>Values</b>     | 0 to 255                                                                                          |

## managed-configuration

|                    |                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] managed-configuration</b>                                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>rtr-adv                                                                                                                                                                                                                |
| <b>Description</b> | This command sets the managed address configuration flag. This flag indicates that DHCPv6 is available for address configuration in addition to any address auto-configured using stateless address auto-configuration. See RFC 3315 for additional details. |
| <b>Default</b>     | no managed-configuration                                                                                                                                                                                                                                     |

## max-advertisement

|                    |                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>max-advertisement <i>seconds</i></b><br><b>no max-advertisement</b>                                    |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>rtr-adv<br>config>service>ies>sub-if>ipv6>rtr-adv                   |
| <b>Description</b> | This command configures the maximum interval between sending router advertisement messages.               |
| <b>Default</b>     | max-advertisement 900                                                                                     |
| <b>Parameters</b>  | <i>seconds</i> — Specifies the maximum interval in seconds between sending router advertisement messages. |
| <b>Values</b>      | 900 to 1800                                                                                               |

## min-advertisement

|                    |                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>min-advertisement <i>seconds</i></b><br><b>no min-advertisement</b>                      |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>rtr-adv<br>config>service>ies>sub-if>ipv6>rtr-adv     |
| <b>Description</b> | This command configures the minimum interval between sending router advertisement messages. |
| <b>Default</b>     | min-advertisement 900                                                                       |

---

|                   |                                                                                                           |
|-------------------|-----------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>seconds</i> — Specifies the minimum interval in seconds between sending router advertisement messages. |
| <b>Values</b>     | 900 to 1350                                                                                               |

## mtu

|                    |                                                                                    |
|--------------------|------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mtu bytes</b><br><b>no mtu</b>                                                  |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>rtr-adv                                      |
| <b>Description</b> | This command configures the MTU for the nodes to use to send packets on the link.  |
| <b>Default</b>     | no mtu                                                                             |
| <b>Parameters</b>  | <i>bytes</i> — Specifies the MTU for the nodes to use to send packets on the link. |
| <b>Values</b>      | 1280 to 9212                                                                       |

## other-stateful-configuration

|                    |                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] other-stateful-configuration</b>                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>rtr-adv                                                                                                                                                                                                                                                                                |
| <b>Description</b> | This command sets the "other configuration" flag. This flag indicates that DHCPv6 is available for autoconfiguration of other (non-address) information such as DNS-related information or information on other servers in the network. See RFC 3736, <i>Stateless Dynamic Host Configuration Protocol (DHCP) for IPv6</i> . |
| <b>Default</b>     | no other-stateful-configuration                                                                                                                                                                                                                                                                                              |

## prefix-options

|                    |                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] prefix-options</b>                                                                                                                                                  |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>rtr-adv                                                                                                                               |
| <b>Description</b> | This command configures Router Advertisement parameters for IPv6 prefixes returned via RADIUS Framed-IPv6-Prefix. All prefixes will inherit these configuration parameters. |
| <b>Default</b>     | no prefix-options                                                                                                                                                           |

---

## autonomous

|                    |                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] autonomous</b>                                                                     |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>rtr-adv>pfx-opt                                      |
| <b>Description</b> | This command specifies whether the prefix can be used for stateless address configuration. |
| <b>Default</b>     | no autonomous                                                                              |

## preferred-lifetime

|                    |                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>preferred-lifetime</b> [ <i>seconds</i>   <i>infinite</i> ]<br><b>no preferred-lifetime</b>                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>rtr-adv>pfx-opt                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | This command configures the remaining length of time in seconds that this prefix will continue to be preferred, for example, time until deprecation. The address generated from a deprecated prefix should not be used as a source address in new communications, but packets received on such an interface are processed as expected. |
| <b>Default</b>     | preferred-lifetime 3600                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <i>seconds</i> — Specifies a decimal time interval in seconds.<br><b>Values</b> 0 to 4294967295<br><i>infinite</i> — Specifies a 0xffffffff value, Dec = 4294967295.                                                                                                                                                                   |

## valid-lifetime

|                    |                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>valid-lifetime</b> [ <i>seconds</i>   <i>infinite</i> ]<br><b>no valid-lifetime</b>                                                                                                                                                                                                                 |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>rtr-adv>pfx-opt                                                                                                                                                                                                                                                  |
| <b>Description</b> | This command specifies the length of time in seconds that the prefix is valid for the purpose of on-link determination. A value of all one bits (0xffffffff) represents infinity. The address generated from an invalidated prefix should not appear as the destination or source address of a packet. |
| <b>Default</b>     | valid-lifetime 86400                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <i>seconds</i> — Specifies a decimal time interval in seconds.<br><b>Values</b> 0 to 4294967295<br><i>infinite</i> — Specifies a 0xffffffff value, Dec = 4294967295.                                                                                                                                   |



---

## reachable-time

|                    |                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>reachable-time</b> <i>milliseconds</i><br><b>no reachable-time</b>                                                                               |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>rtr-adv                                                                                                       |
| <b>Description</b> | This command configures how long this router should be considered reachable by other nodes on the link after receiving a reachability confirmation. |
| <b>Default</b>     | no reachable-time                                                                                                                                   |
| <b>Parameters</b>  | <i>milliseconds</i> — The length of time the router should be considered reachable for default router selection.<br><b>Values</b> 0 to 3600000      |

## retransmit-time

|                    |                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>retransmit-time</b> <i>milliseconds</i><br><b>no retransmit-time</b>                        |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>rtr-adv                                                  |
| <b>Description</b> | This command configures the retransmission frequency of neighbor solicitation messages.        |
| <b>Default</b>     | no retransmit-time                                                                             |
| <b>Parameters</b>  | <i>milliseconds</i> — Specifies how often retransmissions occur.<br><b>Values</b> 0 to 1800000 |

## router-lifetime

|                    |                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>router-lifetime</b> <i>seconds</i><br><b>router-lifetime no-default-router</b><br><b>no router-lifetime</b>                                                                                                     |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>rtr-adv                                                                                                                                                                      |
| <b>Description</b> | This command sets the router lifetime. A value of zero indicates this router should not be used by hosts as a default router.                                                                                      |
| <b>Default</b>     | router-lifetime 4500                                                                                                                                                                                               |
| <b>Parameters</b>  | <i>seconds</i> — Specifies how long the router is valid for default router selection.<br><b>Values</b> 2700 to 9000<br><b>no-default-router</b> — Indicates that the router is not to be used as a default router. |

---

## dhcp6

|                    |                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] dhcp6</b>                                                                                                                                  |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6                                                                                                              |
| <b>Description</b> | This command allows access to the DHCP6 context within the group interface configuration. Within this context, DHCP6 parameters can be configured. |
| <b>Default</b>     | no dhcp6                                                                                                                                           |

## proxy-server

|                    |                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] proxy-server</b>                                                                                                                                  |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>dhcp6                                                                                                               |
| <b>Description</b> | This command allows access to the DHCP6 proxy server context. Within this context, DHCP6 proxy server parameters of the group interface can be configured |
| <b>Default</b>     | no proxy-server                                                                                                                                           |

## client-applications

|                    |                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>client-applications [dhcp] [ppp]</b><br><b>no client-applications</b>                                      |
| <b>Context</b>     | config>services>ies>sub-if>grp-if>ipv6>dhcp6>proxy-server                                                     |
| <b>Description</b> | This command configures the client host types to which the DHCP6 proxy server is allowed to assign addresses. |
| <b>Parameters</b>  | <b>dhcp</b> — Specifies IP over Ethernet hosts.<br><b>ppp</b> — Specifies PPP over Ethernet hosts.            |

## preferred-lifetime

|                    |                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>preferred-lifetime infinite</b><br><b>preferred-lifetime [days days] [hrs hours] [min minutes] [sec seconds]</b><br><b>no preferred-lifetime</b>                            |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy-server                                                                                                                       |
| <b>Description</b> | The preferred lifetime for the IPv6 prefix or address in the option, expressed in units of seconds. When the preferred lifetime expires, any derived addresses are deprecated. |

---

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | preferred-lifetime hrs 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b> | <p><b>infinite</b> — Specifies that the preferred lifetime is infinite.</p> <p><b>days</b> <i>days</i> — Specifies the number of days of a preferred lifetime.</p> <p><b>Values</b> 0 to 49710</p> <p><b>hrs</b> <i>hours</i> — Specifies the number of hours of a preferred lifetime.</p> <p><b>Values</b> 0 to 23</p> <p><b>min</b> <i>minutes</i> — Specifies the number of minutes of a preferred lifetime.</p> <p><b>Values</b> 0 to 59</p> <p><b>sec</b> <i>seconds</i> — Specifies the number of seconds of a preferred lifetime.</p> <p><b>Values</b> 0 to 59</p> |

## rebind-timer

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>rebind-timer</b> [ <b>days</b> <i>days</i> ] [ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ]<br><b>no rebind-timer</b>                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy-server                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | This command configures the rebind-timer (T2), the time at which the client contacts any available server to extend the lifetimes of the addresses or prefixes assigned to the client.                                                                                                                                                                                                                                                                             |
| <b>Default</b>     | rebind-timer min 48                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><b>days</b> <i>days</i> — Specifies the number of days of a rebind timer.</p> <p><b>Values</b> 0 to 14</p> <p><b>hrs</b> <i>hours</i> — Specifies the number of hours of a rebind timer.</p> <p><b>Values</b> 0 to 23</p> <p><b>min</b> <i>minutes</i> — Specifies the number of minutes of a rebind timer.</p> <p><b>Values</b> 0 to 59</p> <p><b>sec</b> <i>seconds</i> — Specifies the number of seconds of a rebind timer.</p> <p><b>Values</b> 0 to 59</p> |

## renew-timer

|                |                                                                                                                                                                 |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>renew-timer</b> [ <b>days</b> <i>days</i> ] [ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ]<br><b>no renew-timer</b> |
| <b>Context</b> | config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy-server                                                                                                        |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command configures the renew-timer (T1), the time at which the client contacts the server from which the addresses in the IA_NA or IA_PD were obtained to extend the lifetimes of the addresses or prefixes assigned to the client.                                                                                                                                                                                                                      |
| <b>Default</b>     | renew-timer min 30                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <p><b>days</b> <i>days</i> — Specifies the number of days of a renew timer.</p> <p><b>Values</b> 0 to 7</p> <p><b>hrs</b> <i>hours</i> — Specifies the number of hours of a renew timer.</p> <p><b>Values</b> 0 to 23</p> <p><b>min</b> <i>minutes</i> — Specifies the number of minutes of a renew timer.</p> <p><b>Values</b> 0 to 59</p> <p><b>sec</b> <i>seconds</i> — Specifies the number of seconds of a renew timer.</p> <p><b>Values</b> 0 to 59</p> |

## valid-lifetime

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>valid-lifetime infinite</b><br><b>valid-lifetime</b> [ <b>days</b> <i>days</i> ] [ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ]<br><b>no valid-lifetime</b>                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy-server                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | The valid lifetime for the IPv6 prefix or address in the option, expressed in units of seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>     | valid-lifetime days 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <p><b>infinite</b> — Specifies that the valid lifetime is infinite.</p> <p><b>days</b> <i>days</i> — Specifies the number of days of a valid lifetime.</p> <p><b>Values</b> 0 to 49710</p> <p><b>hrs</b> <i>hours</i> — Specifies the number of hours of a valid lifetime.</p> <p><b>Values</b> 0 to 23</p> <p><b>min</b> <i>minutes</i> — Specifies the number of minutes of a valid lifetime.</p> <p><b>Values</b> 0 to 59</p> <p><b>sec</b> <i>seconds</i> — Specifies the number of seconds of a valid lifetime.</p> <p><b>Values</b> 0 to 59</p> |

## load-balancing

|                    |                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>load-balancing</b>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>service>ies>if                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | This command enables the load-balancing context to configure interface per-flow load balancing options that will apply to traffic entering this interface and egressing over a LAG/ECMP on system-egress. This is a per interface setting. For load balancing options that can also be enabled on the system level, the options enabled on the interface level overwrite system level configurations. |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                                                                                                                                                                   |

## egr-ip-load-balancing

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>egr-ip-load-balancing {source   destination   inner-ip}</b><br><b>no egr-ip-load-balancing</b>                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>service>ies>if>load-balancing                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | This command specifies whether to include the source address or destination address or both in the LAG/ECMP hash on IP interfaces. Additionally, when l4-load-balancing is enabled, the command also applies to the inclusion of source/destination port in the hash inputs.<br><br>The <b>no</b> form of this command includes both source and destination parameters.                                                                                                                              |
| <b>Default</b>     | no egr-ip-load-balancing                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <i>source</i> — Specifies using the source address and, if l4-load balancing is enabled, the source port in the hash, ignore destination address/port.<br><br><i>destination</i> — Specifies using the destination address and, if l4-load balancing is enabled, the destination port in the hash, ignore source address/port.<br><br><i>inner-ip</i> — Specifies using the inner IP header parameters instead of the outer IP header parameters in the LAG/ECMP hash for IPv4 encapsulated traffic. |

## spi-load-balancing

|                    |                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] spi-load-balancing</b>                                                                                                                                             |
| <b>Context</b>     | config>service>ies>if>load-balancing                                                                                                                                       |
| <b>Description</b> | This command enables use of the SPI in hashing for ESP/AH encrypted IPv4/v6 traffic. This is a per interface setting.<br><br>The <b>no</b> form disables the SPI function. |

---

**Default** no spi-load-balancing

## teid-load-balancing

**Syntax** [no] teid-load-balancing

**Context** config>service>ies>if>load-balancing

**Description** This command enables inclusion of TEID in hashing for GTP-U/C encapsulates traffic for GTPv1/GTPv2. The **no** form of this command ignores TEID in hashing.

**Default** no teid-load-balancing

## local-dhcp-server

**Syntax** local-dhcp-server *local-server-name*  
no local-dhcp-server

**Context** config>service>ies>if

**Description** This command assigns a DHCP server to the interface.

**Parameters** *local-server-name* — Specifies an existing local server name.

## local-proxy-arp

**Syntax** [no] local-proxy-arp

**Context** config>service>ies>if  
config>service>ies>sub-if>grp-if

**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 thus will become the forwarding point for all traffic between hosts in that subnet. When local-proxy-arp is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

**Default** ies>if: no local-proxy-arp  
ies>sub-if>grp-if: local-proxy-arp (7750 SR)

## loopback

**Syntax** [no] loopback

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>ies>if                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>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.</p> <p>You can configure an IES interface as a loopback interface by issuing the <b>loopback</b> command instead of the <b>sap</b> 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.</p> |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## ip-helper-address

|                    |                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ip-helper-address</b> <i>gateway-address</i><br><b>no ip-helper-address</b>                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>service>ies>if                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command enables broadcast UDP packets received on the associated interface to be redirected to the specified gateway address and then forwarded on to the gateway.</p> <p>The <b>no</b> form of this command removes the gateway address from the interface configuration and stops the UDP broadcast redirect function.</p> |
| <b>Parameters</b>  | <i>gateway-address</i> — Specifies the IPv4 address of the target UDP broadcast gateway.                                                                                                                                                                                                                                             |

## ip-mtu

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ip-mtu</b> <i>octets</i><br><b>no ip-mtu</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>service>ies>if<br>config>service>ies>if>sap>ip-tunnel<br>config>service>ies>subscriber-interface                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command configures the IP maximum transmit unit (packet) for this interface.</p> <p>Because this connects a Layer 2 to a Layer 3 service, this parameter can be adjusted under the IES interface.</p> <p>The MTU that is advertised from the IES size is:</p> <p>MINIMUM((SdpOperPathMtu - EtherHeaderSize), (Configured ip-mtu))</p> <p>By default (for Ethernet network interface) if no ip-mtu is configured it is (1568 - 14) = 1554.</p> <p>The <b>no</b> form of the command returns the default value.</p> |

---

**Default**    no ip-mtu

## reassemble

**Syntax**    **reassemble** [*wait-msecs*]  
              **no reassemble**

**Context**    config>service>ies>if>sap>ip-tunnel

**Description**    This command configures the maximum number of seconds to wait to receive all fragments of a particular IPSec or GRE packet for reassembly.

**Default**    no reassemble

**Parameters**    *wait-msecs* — Specifies the reassembly wait time.

**Values**    1 to 5000 ms in 100 increments

## lag-per-link-hash

**Syntax**    **lag-per-link-hash class** {1 | 2 | 3} **weight** *weight*  
              **no lag-per-link-hash**

**Context**    config>service>ies>if>sap  
              config>service>ies>sub-if>grp-if>sap

**Description**    This command configures weight and class to this SAP to be used on LAG egress when the LAG uses weighted per-link-hash.

              The **no** form of this command restores default configuration.

**Default**    no lag-per-link-hash (equivalent to weight 1 class 1)

**Parameters**    *weight* — Specifies the weight.

**Values**    1 to 1024

## mac

**Syntax**    **mac** *ieee-address*  
              **no mac** [*ieee-address*]

**Context**    config>service>ies>if  
              config>service>ies>sub-if>grp-if

**Description**    This command assigns a specific MAC address to an IES IP interface.



For Routed Central Office (CO), a group interface has no IP address explicitly configured but inherits an address from the parent subscriber interface when needed. For example, a MAC will respond to an ARP request when an ARP is requested for one of the IPs associated with the subscriber interface through the group interface.

The **no** form of the command returns the MAC address of the IP interface to the default value.

- Default** the physical MAC address associated with the Ethernet interface that the SAP is configured on (the default MAC address assigned to the interface, assigned by the system)
- Parameters** *ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-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.

## monitor-oper-group

- Syntax** **monitor-oper-group** *name*  
**no monitor-oper-group**
- Context** config>service>ies>if
- Description** This command specifies the operational group to be monitored by the object under which it is configured. The oper-group name must be already configured under the config>service context before its name is referenced in this command.
- The **no** form of the command removes the association from the configuration.
- Default** no monitor-oper-group
- Parameters** *name* — Specifies a character string of maximum 32 ASCII characters identifying the group instance.

## multicast-network-domain

- Syntax** **multicast-network-domain** *multicast-network-domain*  
**no multicast-network-domain**
- Context** config>service>ies>if
- Description** This command is used to enable efficient multicast replication over a spoke SDP. Multicast traffic is copied to only a subset of network interfaces that may be used as egress for a spoke SDP. A network domain is defined by associating multiple interfaces to a logical group that may participate in multicast replication for a spoke SDP.
- The **no** form of command disables efficient multicast replication to a network domain for a spoke SDP and traffic is replicated to all forwarding complexes.

---

**Default** no multicast-network-domain

## secondary

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>secondary</b> { <i>ip-address/mask</i>   <i>ip-address netmask</i> } [ <b>broadcast all-ones</b>   <b>host-ones</b> ] [ <b>igmp-inhibit</b> ] [ <b>track-srrp</b> <i>srrp-instance</i> ]<br><b>no secondary</b> { <i>ip-address/mask</i>   <i>ip-address netmask</i> }                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies>if                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | This command assigns a secondary IP address/IP subnet/broadcast address format to the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <p><i>ip-address</i> — The IP address of the IP interface. The <i>ip-address</i> portion of the <b>address</b> command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><i>mask</i> — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the <i>ip-address</i> from a traditional dotted decimal mask. The <i>mask</i> parameter indicates the complete mask that will be used in a logical 'AND' function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. A mask of 255.255.255.255 is reserved for system IP addresses.</p> <p><i>netmask</i> — 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.</p> <p><b>broadcast</b> — The optional <b>broadcast</b> parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is <b>host-ones</b> which indicates a subnet broadcast address. Use this parameter to change the broadcast address to <b>all-ones</b> or revert back to a broadcast address of <b>host-ones</b>. The broadcast format on an IP interface can be specified when the IP address is assigned or changed.</p> <p>This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (<b>all-ones</b>) or the valid subnet broadcast address (<b>host-ones</b>) will be received by the IP interface. (<i>Default: host-ones</i>)</p> <p><b>all-ones</b> — The <b>all-ones</b> keyword following the <b>broadcast</b> parameter specifies the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.</p> |

**host-ones** — The **host-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default broadcast address used by an IP interface.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

**igp-inhibit** — The optional **igp-inhibit** parameter signals that the given secondary IP interface should not be recognized as a local interface by the running IGP. For OSPF and IS-IS, this means that the specified secondary IP interfaces will not be injected and used as passive interfaces and will not be advertised as internal IP interfaces into the IGP's link state database. For RIP, this means that these secondary IP interfaces will not source RIP updates.

**track-srrp** *srrp-instance* — Specifies the SRRP instance ID that this interface route needs to track.

**Values** 1 to 4294967295

## shcv-policy-ipv4

|                    |                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>shcv-policy-ipv4</b> <i>policy-name</i><br><b>no shcv-policy-ipv4</b>                          |
| <b>Context</b>     | config>service>ies>if<br>config>service>ies>sub-if>grp-if                                         |
| <b>Description</b> | This command specifies the Subscriber Host Connectivity Verification (SHCV) policy for IPv4 only. |

The **no** form of the command removes the policy name from the SAP configuration.

## shcv-policy-ipv6

|                    |                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>shcv-policy-ipv6</b> <i>policy-name</i><br><b>no shcv-policy-ipv6</b>                          |
| <b>Context</b>     | config>service>ies>if<br>config>service>ies>sub-if>grp-if                                         |
| <b>Description</b> | This command specifies the Subscriber Host Connectivity Verification (SHCV) policy for IPv6 only. |

The **no** form of the command removes the policy name from the SAP configuration.

---

## static-arp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>static-arp</b> <i>ieee-mac-address</i> <b>unnumbered</b><br><b>static-arp</b> <i>ip-address</i> <i>ieee-mac-address</i><br><b>no static-arp</b> [ <i>ieee-mac-address</i> ] <b>unnumbered</b><br><b>no static-arp</b> <i>ip-address</i> [ <i>ieee-mac-address</i> ]                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>service>ies>if                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP appears in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.</p> <p>If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.</p> <p>The <b>no</b> form of the command removes a static ARP entry.</p>                    |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <p><i>ip-address</i> — Specifies the IP address for the static ARP in IP address dotted decimal notation.</p> <p><i>ieee-mac-address</i> — 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.</p> <p><b>unnumbered</b> — Specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. Once this command is configured, it overrides any dynamic ARP.</p> |

## static-tunnel-redundant-next-hop

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>static-tunnel-redundant-next-hop</b> <i>ip-address</i><br><b>no static-tunnel-redundant-next-hop</b>                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>if                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>This command specifies redundant next-hop address on public or private IPsec interface (with public or private tunnel-sap) for static IPsec tunnel. The specified next-hop address will be used by standby node to shunt traffic to master in case of it receives them.</p> <p>The next-hop address will be resolved in routing table of corresponding service.</p> <p>The no form of the command removes the address from the interface configuration.</p> |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

---

**Parameters**    *ip-address* — Specifies the static ISA tunnel redundant next-hop address.

## tos-marking-state

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>tos-marking-state {trusted   untrusted}</b><br><b>no tos-marking-state</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>service>ies>if<br>config>service>ies>sub-if>grp-if                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command is used to change the default trusted state to a non-trusted state. When unset or reverted to the trusted default, the ToS field will not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set, in which case the egress network interface treats all IES and network IP interface as untrusted.</p> <p>When the ingress interface is set to untrusted, all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface. The egress network remarking rules also apply to the ToS field of IP packets routed using IGP shortcuts (tunneled to a remote next-hop). However, the tunnel QoS markings are always derived from the egress network QoS definitions.</p> <p>Egress marking and remarking is based on the internal forwarding class and profile state of the packet once it reaches the egress interface. The forwarding class is derived from ingress classification functions. The profile of a packet is either derived from ingress classification or ingress policing.</p> <p>The default marking state for network IP interfaces is trusted. This is equivalent to declaring no tos-marking-state on the network IP interface. When undefined or set to tos-marking-state trusted, the trusted state of the interface will not be displayed when using show config or show info unless the detail parameter is given. The <b>save config</b> command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified.</p> <p>The <b>no</b> form of the command is used to restore the trusted state to a network IP interface. This is equivalent to executing the tos-marking-state trusted command.</p> |
| <b>Default</b>     | tos-marking-state untrusted                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><b>trusted</b> — The default prevents the ToS field to not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set.</p> <p><b>untrusted</b> — Specifies that all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## unnumbered

**Syntax**    **unnumbered** [{*ip-int-name* | *ip-address*}]

**no unnumbered**

|                    |                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>ies>if<br>config>service>ies>subscriber-interface                                                                                                                                                                                                                                               |
| <b>Description</b> | This command configures the interface as an unnumbered interface. Unnumbered IP interfaces are supported on a SONET/SDH access port with the PPP, ATM, Frame Relay, cisco-HDLC encapsulation. It is not supported on access ports that do not carry IP traffic, but are used for native TDM circuit emulation. |
| <b>Parameters</b>  | <i>ip-int-name</i> — Specifies the name of an IP interface. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.<br><br><i>ip-address</i> — Specifies an IP address.                                                                 |

## wpp

|                    |                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] wpp</b>                                                                                                                                                             |
| <b>Context</b>     | config>service>ies                                                                                                                                                          |
| <b>Description</b> | This command enters the configuration context of web portal protocol (WPP) under router or vprn.<br><br>The <b>no</b> form of this command removes configuration under WPP. |
| <b>Default</b>     | no wpp                                                                                                                                                                      |

## initial-app-profile

|                    |                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>initial-app-profile</b> <i>profile-name</i><br><b>no initial-app-profile</b>                                                                                                   |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>wpp                                                                                                                                              |
| <b>Description</b> | This command specifies the initial app-profile for the hosts created on the group-interface. This initial app-profile is replaced after hosts pass the web portal authentication. |
| <b>Default</b>     | no initial-app-profile                                                                                                                                                            |
| <b>Parameters</b>  | <i>profile-name</i> — Specifies the name of app-profile.                                                                                                                          |

## initial-sla-profile

|               |                                                                                 |
|---------------|---------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>initial-sla-profile</b> <i>profile-name</i><br><b>no initial-sla-profile</b> |
|---------------|---------------------------------------------------------------------------------|

---

|                    |                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>wpp                                                                                                                                                                 |
| <b>Description</b> | This command specifies the initial sla-profile for the hosts created on the group-interface. This initial sla-profile is replaced after hosts pass the web portal authentication. |
| <b>Default</b>     | no initial-sla-profile                                                                                                                                                            |
| <b>Parameters</b>  | <i>profile-name</i> — Specifies the name of sla-profile.                                                                                                                          |

## initial-sub-profile

|                    |                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>initial-sub-profile</b> <i>profile-name</i><br><b>no initial-sub-profile</b>                                                                                                    |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>wpp                                                                                                                                               |
| <b>Description</b> | This command specifies the initial sub-profile for the hosts created on the group-interface. This initial sub-profile will be replaced after hosts pass web portal authentication. |
| <b>Default</b>     | no initial-sub-profile                                                                                                                                                             |
| <b>Parameters</b>  | <i>profile-name</i> — Specifies the name of sub-profile.                                                                                                                           |

## portal

|                    |                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>portal router</b> <i>router-instance name wpp-portal-name</i><br><b>no portal</b>                                                                             |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>wpp                                                                                                                             |
| <b>Description</b> | This command specifies the web portal server that system talks to for the hosts on the group-interface.                                                          |
| <b>Default</b>     | no portal                                                                                                                                                        |
| <b>Parameters</b>  | <i>router-instance</i> — Specifies the routing-instance that web portal server is defined.<br><i>profile-name</i> — Specifies the name of the web portal server. |

## restore-disconnected

|                    |                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] restore-disconnected</b>                                                                                                                                                                       |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>wpp                                                                                                                                                                   |
| <b>Description</b> | This command enable the behavior that system will restore the <b>initial-sla-profile</b> , <b>initial-sub-profile</b> , or <b>initial-app-profile</b> when hosts disconnects instead of removing them. |

The **no** form of the command specifies that the initial profiles will not be restored after a DHCP host has disconnected.

**Default** restore-disconnected

## urpf-check

**Syntax** [no] **urpf-check**

**Context** config>service>ies>if  
config>service>ies>if>ipv6  
config>service>ies>sub-if>grp-if>ipv6

**Description** This command enables the unicast RPF (uRPF) Check on this interface.

The **no** form of the command disables the uRPF Check on this interface.

**Default** no urpf-check

## vas-if-type

**Syntax** **vas-if-type** {**to-from-access** | **to-from-network** | **to-from-both**}  
**no vas-if-type**

**Context** config>service>ies>if

**Description** This command configures the type of a Value Added Service (VAS) facing interface. To change the **vas-if-type**, the **shutdown** command is required. The **vas-if-type** and **loopback** commands are mutually exclusive.

The **no** form of the command removes the VAS interface type configuration.

**Default** no vas-if-type

**Parameters** **to-from-access** — Used when two separate (**to-from-access** and **to-from-network**) interfaces are used for VAS connectivity. For service chaining, traffic arriving from access interfaces (upstream) is redirected to a PBR target reachable over this interface for upstream VAS processing. Downstream traffic after VAS processing must arrive on this interface, so that the traffic is subject to regular routing but is not subject to AA divert, nor egress subscriber PBR.

**to-from-network** — Used when two separate (**to-from-access** and **to-from-network**) interfaces are used for VAS connectivity. For service chaining, traffic arriving from network interfaces (downstream) is redirected to a PBR target reachable over this interface for downstream VAS processing. Upstream traffic after VAS processing must arrive on this interface, so that regular routing can be applied.



**to-from-both** — Used when a single interface is used for VAS connectivity (no local-to-local traffic). For service chaining, both traffic arriving from access interfaces and from network interfaces is redirected to a PBR target reachable over this interface for upstream/downstream VAS processing. Traffic after VAS processing must arrive on this interface, so that the traffic is subject to regular routing but is not subject to AA divert, nor egress subscriber PBR.

## mode

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mode</b> { <b>strict</b>   <b>loose</b>   <b>strict-no-ecmp</b> }                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>service>ies>if>urpf-check<br>config>service>ies>sub-if>grp-if>ipv6>urpf-check                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | This command specifies the mode of unicast RPF check.<br><br>The <b>no</b> form of the command reverts to the default (strict) mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>     | mode strict                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <b>strict</b> — When specified, uRPF checks whether incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix.<br><br><b>loose</b> — In <b>loose</b> mode, uRPF checks whether incoming packet has source address with a corresponding prefix in the routing table. However, the loose mode does not check whether the interface expects to receive a packet with a specific source address prefix. This object is valid only when <b>urpf-check</b> is enabled.<br><br><b>strict-no-ecmp</b> — When a packet is received on an interface in this mode and the SA matches an ECMP route the packet is dropped by uRPF. |

## vpls

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>vpls</b> <i>service-name</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>service<br>config>service>ies>if                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | The <b>vpls</b> command, within the IP interface context, is used to bind the IP interface to the specified service name (VPLS or I-VPLS).<br><br>The system does not attempt to resolve the service name provided until the IP interface is placed into the administratively up state ( <b>no shutdown</b> ). Once the IP interface is administratively up, the system will scan the available VPLS services that have the <b>allow-ip-int-bind</b> flag set for a VPLS service associated with the name. If the service name is bound to the service name when the IP interface is already in the administratively up state, the system will immediately attempt to resolve the given name. |

If a VPLS service is found associated with the name and with the **allow-ip-int-bind** flag set, the IP interface will be attached to the VPLS service allowing routing to and from the service virtual ports once the IP interface is operational.

A VPLS service associated with the specified name that does not have the **allow-ip-int-bind** 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-bind** flag is set at the time the name is applied, the VPLS service will be 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-bind** flag set, the system will 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-bind** flag set will be attached to the VPLS service. Only one IP interface is allowed to attach to a VPLS service context. No error is generated for the remaining non-attached IP interfaces using the service name.

Once an IP interface is attached to a VPLS service, the name associated with the service cannot be removed or changed until the IP interface name binding is removed. Also, the **allow-ip-int-bind** 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.

### VPRN Hardware Dependency

When a service name is bound to a VPRN IP interface, all SAPs associated with the VPRN service must be on hardware based on the FlexPath2 forwarding plane. Currently, these include the IOM3-XP, the various IMM modules and the SR7710c12. If any SAPs are associated with the wrong hardware type, the service name binding to the VPRN IP interface fails. Once an IP interface within the VPRN service is bound to a service name, attempting to create a SAP on excluded hardware fails.

### IP Interface MTU and Fragmentation

A 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 and CRC) that may be transmitted out a port. The VPLS itself has a service level MTU that defines the largest packet supported by the service. This MTU does not include the local encapsulation overhead for each port (QinQ, Dot1Q, TopQ or SDP service delineation fields and headers) but does include the remainder of the packet. As virtual ports are created in the system, the virtual port cannot become operational unless the configured port MTU minus the virtual port service delineation overhead is greater than or equal to the configured VPLS service MTU. Thus, an operational virtual port 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.

IP interfaces have a configurable up MTU that defines the largest packet that may egress the IP interface without being fragmented. This MTU encompasses the IP portion of the packet and does not include any of the egress DLC header or CRC. This MTU does not affect the size of the largest ingress packet on the IP interface. If the egress IP portion of the packet is larger than the IP interface MTU and the IP header do not fragment flag is not set, the packet is fragmented into smaller packets that will not exceed the configured MTU size. If the do not fragment bit is set, the packet is silently discarded at egress when it exceeds the IP MTU.

When the IP interface is bound to a VPLS service, the IP MTU must be at least 18 bytes less than the VPLS service MTU. This allows for the addition of the minimal Ethernet encapsulation overhead; 6 bytes for the DA, 6 bytes for the SA, 2 bytes for the Etype and 4 bytes for the trailing CRC. Any remaining egress virtual port overhead (Dot1P, Dot1Q, QinQ, TopQ or SDP) required above the minimum is known to be less than the egress ports MTU since the virtual port would not be operational otherwise.

If the IP interface IP MTU value is too large based on the VPLS service MTU, the IP interface will enter the operationally down state until either the IP MTU is adequately lowered or the VPLS service MTU is sufficiently increased.

The **no** form of the command on the IP interface is used to remove the service name binding from the IP interface. If the service name has been resolved to a VPLS service context and the IP interface has been attached to the VPLS service, the IP interface will also be detached from the VPLS service.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | n/a                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b> | <i>service-name</i> — The service-name parameter is required when using the IP interface vpls command and specifies the service name that the system will attempt to resolve to an <b>allow-ip-int-bind</b> 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

|                    |                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ingress</b>                                                                                                                |
| <b>Context</b>     | config>service>ies>if>vpls                                                                                                    |
| <b>Description</b> | The ingress node in this context under the vpls binding is used to define the routed IPv4 and IPv6 optional filter overrides. |

## v4-routed-override-filter

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>v4-routed-override-filter</b> <i>ip-filter-id</i><br><b>no v4-routed-override-filter</b>                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>service>ies>if>vpls>egress                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command configures an IPv4 filter ID that are applied to packets egressing the IES R-VPLS interface. The filter overrides existing egress IPv4 filter applied to VPLS service endpoints such as SAPs or SDPs, if configured.</p> <p>The <b>no</b> form of the command removes the IPv4 routed override filter from the egress IES R-VPLS interface. When removed, egress IPv4 packets will use the IPv4 egress filter applied to the VPLS endpoint, if configured.</p> |
| <b>Parameters</b>  | <i>ip-filter-id</i> — Specifies the IP filter ID. This parameter is required when executing the <b>v4-routed-override-filter</b> command. The specified filter ID must exist as an IPv4 filter within the system or the override command fails.                                                                                                                                                                                                                                |

## v4-routed-override-filter

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>v4-routed-override-filter</b> <i>ip-filter-id</i><br><b>no v4-routed-override-filter</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>service>ies>if>vpls>ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command configures an IPv4 filter ID that is applied to all ingress packets entering the VPLS or I-VPLS service. The filter overrides any existing ingress IPv4 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 IPv4 routed packets use any existing ingress IPv4 filter on the VPLS virtual port.</p> <p>The <b>no</b> form of the command removes the IPv4 routed override filter from the ingress IP interface. When removed, the IPv4 ingress routed packets within a VPLS service attached to the IP interface use the IPv4 ingress filter applied to the packets virtual port, when defined.</p> |
| <b>Parameters</b>  | <i>ip-filter-id</i> — Specifies the IP filter ID. This parameter is required when executing the <b>v4-routed-override-filter</b> command. The specified filter ID must exist as an IPv4 filter within the system or the override command fails.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## v6-routed-override-filter

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>v6-routed-override-filter</b> <i>ipv6-filter-id</i><br><b>no v6-routed-override-filter</b>                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>service>ies>if>vpls>egress                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command configures an IPv6 filter ID that is applied to packets egressing the IES R-VPLS interface. The filter overrides existing egress IPv6 filter applied to VPLS service endpoints such as SAPs or SDPs, if configured.</p> <p>The <b>no</b> form of the command removes the IPv4 routed override filter from the egress IES R-VPLS interface. When removed, egress IPv6 routed packets uses the IPv6 egress filter applied to VPLS endpoint, if configured</p> |
| <b>Parameters</b>  | <i>ipv6-filter-id</i> — Specifies the IPv6 filter ID. This parameter is required when executing the <b>v6-routed-override-filter</b> command. The specified filter ID must exist as an IPv6 filter within the system or the override command fails.                                                                                                                                                                                                                         |

## v6-routed-override-filter

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>v6-routed-override-filter</b> <i>ipv6-filter-id</i><br><b>no v6-routed-override-filter</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>service>ies>if>vpls>ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command configures an IPv6 filter ID that is applied to all ingress packets entering the VPLS or I-VPLS service. The filter overrides any existing ingress IPv6 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 IPv6 routed packets use any existing ingress IPv6 filter on the VPLS virtual port.</p> <p>The no v6-routed-override-filter command is used to remove the IPv6 routed override filter from the ingress IP interface. When removed, the IPv6 ingress routed packets within a VPLS service attached to the IP interface will use the IPv6 ingress filter applied to the packet's virtual port, when defined.</p> |
| <b>Parameters</b>  | <i>ipv6-filter-id</i> — Specifies the IPv6 filter ID. This parameter is required when executing the <b>v6-routed-override-filter</b> command. The specified filter ID must exist as an IPv6 filter within the system or the override command fails.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## egress

|                |                            |
|----------------|----------------------------|
| <b>Syntax</b>  | <b>egress</b>              |
| <b>Context</b> | config>service>ies>if>vpls |

---

|                    |                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | The egress node under the vpls binding is used to define the optional sap-egress QoS policy that will be used for reclassifying the egress forwarding class or profile for routed packets associated with the IP interface on the attached VPLS or I-VPLS service context. |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## reclassify-using-qos

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>reclassify-using-qos</b> <i>policy-id</i><br><b>no reclassify-using-qos</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>service>ies>if>vpls>egress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>The reclassify-using-qos command is used to specify a sap-egress QoS policy that will be used to reclassify the forwarding class and profile of egress routed packets on the VPLS or I-VPLS service. When routed packets associated with the IP interface egress a VPLS SAP, the reclassification rules within the sap-egress QoS policy applied to the SAP are always ignored (even when reclassify-using-qos is not defined).</p> <p>Any queues or policers defined within the specified QoS policy are ignored and are not created on the VPLS egress SAPs. Instead, the routed packets continue to use the forwarding class mappings, queues and policers from the sap-egress QoS policy applied to the egress VPLS SAP.</p> <p>While the specified sap-egress policy ID is applied to an IP interface it cannot be deleted from the system.</p> <p>The <b>no</b> form of the command removes the sap-egress QoS policy used for reclassification from the egress IP interface. When removed, IP routed packets will not be reclassified on the egress SAPs of the VPLS service attached to the IP interface.</p> |
| <b>Parameters</b>  | <i>policy-id</i> — Specifies the SAP egress QoS policy ID. This parameter is required when executing the reclassify-using-qos command. The specified SAP egress QoS ID must exist within the system or the command fails.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## proxy-arp-policy

|                    |                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>proxy-arp-policy</b> <i>policy-name</i> [ <i>policy-name</i> ]<br><b>no proxy-arp-policy</b>                                                              |
| <b>Context</b>     | config>service>ies>if                                                                                                                                        |
| <b>Description</b> | <p>This command configures up to five proxy ARP policies for the interface.</p> <p>The <b>no</b> form of this command disables the proxy ARP capability.</p> |
| <b>Default</b>     | no proxy-arp-policy                                                                                                                                          |

---

**Parameters**     *policy-name* — The export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. The specified name(s) must already be defined. Up to 5 route policies can be specified in a single statement.

## ptp-hw-assist

**Syntax**     **[no] ptp-hw-assist**

**Context**     config>service>ies>if

**Description**     This command configures the 1588 port based timestamping assist function for the interface. This capability is supported on a specific set of hardware. The command may be blocked if not all hardware has the required level of support.

Only one interface per physical port can have ptp-hw-assist enabled.

no ptp-hw-assist

## qos-route-lookup

**Syntax**     **qos-route-lookup [source | destination]**  
**no qos-route-lookup**

**Context**     config>service>ies>if  
config>service>ies>if>ipv6  
config>service>ies>sub-if>grp-if  
config>service>ies>sub-if>grp-if>ipv6

**Description**     This command enables QoS classification of the ingress IP packets on an interface based on the QoS information associated with routes in the forwarding table.

If the optional **destination** parameter is specified and the destination address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the destination address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network qos policy.

If the optional **source** parameter is specified and the source address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the source address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network qos policy.

If neither the optional **source** or **destination** parameter is present, then the default is **destination** address matching.

The functionality enabled by the qos-route-lookup command can be applied to IPv4 packets or IPv6 packets on an interface, depending on whether it is present at the interface context (applies to IPv4) or the interface>ipv6 context (applies to IPv6). Subscriber management group interfaces also do not support the source QPPB option.

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

|                   |                                                                                                                                                                                                                                                                                        |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | destination                                                                                                                                                                                                                                                                            |
| <b>Parameters</b> | <b>source</b> — Enables QoS classification of incoming IP packets based on the source address matching a route with QoS information.<br><b>destination</b> — Enables QoS classification of incoming IP packets based on the destination address matching a route with QoS information. |

## secure-nd

|                    |                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] secure-nd                                                                                                                                              |
| <b>Context</b>     | config>service>ies>if>ipv6                                                                                                                                  |
| <b>Description</b> | This command enables Secure Neighbor Discovery (SeND) on the IPv6 interface.<br>The <b>no</b> form of the command reverts to the default and disabled SeND. |

## allow-unsecured-msgs

|                    |                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] allow-unsecured-msgs                                                                                                                                                                                                                   |
| <b>Context</b>     | config>service>ies>if>ipv6>secure-nd                                                                                                                                                                                                        |
| <b>Description</b> | This command specifies whether unsecured messages are accepted. When Secure Neighbor Discovery (SeND) is enabled, only secure messages are accepted by default.<br>The <b>no</b> form of the command disables accepting unsecured messages. |

## link-local-modifier

|                    |                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>link-local-modifier</b> <i>modifier</i><br>[no] link-local-modifier                                   |
| <b>Context</b>     | config>service>ies>if>ipv6>secure-nd                                                                     |
| <b>Description</b> | This command configures the Cryptographically Generated Address (CGA) modifier for link-local addresses. |



**Parameters**    modifier — Specifies the modifier in 32 hexadecimal nibbles.  
**Values**        0x0–0xFFFFFFFF

## public-key-min-bits

**Syntax**        **public-key-min-bits** *bits*  
                  [no] **public-key-min-bits**

**Context**        config>service>ies>if>ipv6>secure-nd

**Description**    This command configures the minimum acceptable key length for public keys used in the generation of a Cryptographically Generated Address (CGA).

**Parameters**    *bits* — Specifies the number of bits.  
**Values**        512 to 1024

## security-parameter

**Syntax**        **security-parameter** *sec*  
                  [no] **security-parameter**

**Context**        config>service>ies>if>ipv6>secure-nd

**Description**    This command configures the security parameter used in the generation of a Cryptographically Generated Address (CGA).

**Parameters**    *sec* — Specifies the security parameter.  
**Values**        0 to 1

## shutdown

**Syntax**        [no] **shutdown**

**Context**        config>service>ies>if>ipv6>secure-nd

**Description**    This command enables or disables Secure Neighbor Discovery (SeND) on the interface.

## stale-time

**Syntax**        **stale-time** *seconds*  
                  **no stale-time**

**Context**        config>service>ies>ipv6

---

```
config>service>ies>if>ipv6
```

**Description** This command configures the time a neighbor discovery cache entry can remain stale before being removed.

The **no** form of the command removes the stale-time value.

**Default** no stale-time

**Parameters** *seconds* — The allowed stale time (in seconds) before a neighbor discovery cache entry is removed.

**Values** 60 to 65535

## tcp-mss

**Syntax** **tcp-mss** *mss-value*  
**no tcp-mss**

**Context** config>service>ies>if  
config>service>ies>if>ipv6

**Description** This command statically sets the TCP maximum segment size (MSS) for TCP connections originated from the associated IP interface to the specified value.

The **no** form of the command removes the static value and allows the TCP MSS value to be calculated based on the IP MTU value by subtracting the base IP and TCP header lengths from the IP MTU value ( $\text{tcp\_mss} = \text{ip\_mtu} - 40$ ).

**Default** no tcp-mss

**Parameters** *mss-value* — The TCP MSS value that should be used in the TCP SYN packet during the three-way handshake negotiation of a TCP connection.

**Note:**  $9158 = \text{max-IP\_MTU} (9198) - 40$

**Values** 536 to 9158 (IPv4)  
1220 to 9138 (IPv6)

## remote-proxy-arp

**Syntax** [**no**] **remote-proxy-arp**

**Context** config>service>ies>if  
config>service>ies>sub-if>grp-if

**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.

**Default** no remote-proxy-arp

## ipv6

**Syntax** [no] ipv6

**Context** config>services>ies>sub-if

**Description** This command enables IPv6 forwarding on the specified subscriber-interface.

**Default** no ipv6

## subscriber-prefixes

**Syntax** [no] subscriber-prefixes

**Context** config>services>ies>sub-if>ipv6

**Description** This command specifies aggregate off-link subscriber prefixes associated with this subscriber interface. Individual prefixes are specified under the prefix context list aggregate routes in which the next-hop is indirect via the subscriber interface.

## prefix

**Syntax** **prefix** *ipv6-address/prefix-length* [pd] [wan-host]  
**no prefix** *ipv6-address/prefix-length*

**Context** config>services>ies>sub-if>ipv6>sub-prefixes

**Description** This command allows a list of prefixes (using the prefix command multiple times) to be routed to hosts associated with this subscriber interface. Each prefix will be represented in the associated FIB with a reference to the subscriber interface. Prefixes are defined as being for prefix delegation (pd) or use on a WAN interface or host (wan-host).

**Parameters** *ipv6-address* — Specifies the 128-bit IPv6 address.

**Values** 128-bit hexadecimal IPv6 address in compressed form.

*prefix-length* — Specifies the length of any associated aggregate prefix.

**Values** 32-63

**pd** — Specifies that this aggregate is used by IPv6 ESM hosts for DHCPv6 prefix-delegation.

**wan-host** — Specifies that this aggregate is used by IPv6 ESM hosts for local addressing or by a routing gateway's WAN interface.

## allow-unmatching-prefixes

|                    |                                                                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] allow-unmatching-prefixes</b>                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>service>ies>sub-if>ipv6                                                                                                                                                                                                                                           |
| <b>Description</b> | This command allows address assignment to PPPoX hosts in cases where the assigned address falls outside the range of the configured subnets below the subscriber interface. Alternatively, if the interface is configured as unnumbered, this command cannot be enabled. |
| <b>Default</b>     | no allow-unmatching-prefixes                                                                                                                                                                                                                                             |

## delegated-prefix-length

|                    |                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>delegated-prefix-length</b> <i>prefix-length</i><br><b>delegated-prefix-length</b> <b>variable</b><br><b>nodelegated-prefix-length</b> |
| <b>Context</b>     | config>services>ies>sub-if>ipv6                                                                                                           |
| <b>Description</b> | This command defines the prefix-length used for all DHCPv6 prefix delegations on this subscriber interface.                               |
| <b>Default</b>     | delegated-prefix-length 64                                                                                                                |
| <b>Parameters</b>  | <i>prefix-length</i> — Specifies the prefix length in use on this subscriber interface for DHCPv6 IA_PD.<br><b>Values</b> 48 to 64        |

## redundant-interface

|                    |                                                                                    |
|--------------------|------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>redundant-interface</b> <i>red-ip-int-name</i><br><b>no redundant-interface</b> |
| <b>Context</b>     | config>service>ies<br>config>service>ies>sub-if>grp-if                             |
| <b>Description</b> | This command configures a redundant interface used for dual homing.                |
| <b>Parameters</b>  | <i>red-ip-int-name</i> — Specifies the redundant IP interface name.                |

## arp-host

|                    |                                                                   |
|--------------------|-------------------------------------------------------------------|
| <b>Syntax</b>      | <b>arp-host</b>                                                   |
| <b>Context</b>     | config>service>ies>sub-if>grp-if                                  |
| <b>Description</b> | This command enters the context to configure ARP host parameters. |

## host-limit

|                    |                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>host-limit</b> <i>max-num-hosts</i><br><b>no host-limit</b>                                |
| <b>Context</b>     | config>service>ies>sub-if>grp-if                                                              |
| <b>Description</b> | This command configures the maximum number of ARP hosts.                                      |
| <b>Parameters</b>  | <i>max-num-hosts</i> — Specifies the maximum number of ARP hosts.<br><b>Values</b> 1 to 32767 |

## min-auth-interval

|                    |                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>min-auth-interval</b> <i>min-auth-interval</i><br><b>no min-auth-interval</b>                     |
| <b>Context</b>     | config>service>ies>sub-if>grp-if                                                                     |
| <b>Description</b> | This command configures the minimum authentication interval.                                         |
| <b>Parameters</b>  | <i>min-auth-interval</i> — Specifies the minimum authentication interval.<br><b>Values</b> 1 to 6000 |

## sap-host-limit

|                    |                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sap-host-limit</b> <i>max-num-hosts-sap</i><br><b>no sap-host-limit</b>                                                              |
| <b>Context</b>     | config>service>ies>sub-if>grp-if                                                                                                        |
| <b>Description</b> | This command configures the maximum number of ARP hosts per SAP.                                                                        |
| <b>Parameters</b>  | <i>max-num-hosts-sap</i> — Specifies the maximum number of ARP hosts per SAP allowed on this IES interface.<br><b>Values</b> 1 to 32767 |

---

## frame-relay

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>frame-relay</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>service>ies>if>sap                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command allows access to the context to configure the Frame Relay Local Management Interface (LMI) operational parameters for a SONET/SDH PoS link, a DS-0 channel group, or a DS-3/E-3 port or channel.</p> <p>The port's <b>mode</b> must be set to <b>access</b> in <b>config&gt;port&gt;sonet-sdh&gt;path&gt;mode access</b> context.</p> <p>The port's encapsulation type must be set to <b>frame-relay</b> in the <b>config&gt;port&gt;sonet-sdh&gt;path&gt;encap-type frame-relay</b> context.</p> <p>The <b>no</b> form of this command removes the Frame Relay LMI operational parameters.</p> |

## delivery-service

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>delivery-service</b> <i>service-id</i><br><b>delivery-service name</b> <i>service-name</i><br><b>no delivery-service</b>                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>service>if>ies>sap<br>config>service>if>vprn>sap>ip-tunnel                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command sets the delivery service for encapsulated packets associated with a particular tunnel. This is the IES or VPRN service where the encapsulated packets are injected and terminated. The delivery service may be the same service that owns the private tunnel SAP associated with the tunnel. The tunnel does not come up until a valid delivery service is configured.</p> <p>The <b>no</b> form of the command deletes the delivery-service from the tunnel configuration.</p> |
| <b>Parameters</b>  | <p><i>service-id</i> — Identifies the service used to originate and terminate the encapsulated packets belonging to the tunnel.</p> <p><b>Values</b> 1 to 2147483648</p> <p><i>service-name</i> — Identifies the service used to originate and terminate the encapsulated packets belonging to the GRE tunnel.</p> <p><b>Values</b> 1 to 64 characters</p>                                                                                                                                       |

## dest-ip

|               |                                       |
|---------------|---------------------------------------|
| <b>Syntax</b> | <b>[no] dest-ip</b> <i>ip-address</i> |
|---------------|---------------------------------------|

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                             |  |  |              |              |         |  |              |                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|--|--|--------------|--------------|---------|--|--------------|---------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>ies>if>sap>ip-tunnel<br>config>service>vprn>if>sap>ip-tunnel                                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                             |  |  |              |              |         |  |              |                                                                                             |
| <b>Description</b> | <p>This command configures a private IPv4 or IPv6 address of the remote tunnel endpoint. A tunnel can have up to 16 <b>dest-ip</b> commands. At least one <b>dest-ip</b> address is required in the configuration of a tunnel. A tunnel does not come up operationally unless all <b>dest-ip</b> addresses are reachable (part of a local subnet).</p> <p>Unnumbered interfaces are not supported.</p> <p>The <b>no</b> form of the command deletes the destination IP of the tunnel.</p> |                                                                                             |  |  |              |              |         |  |              |                                                                                             |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                                                                             |  |  |              |              |         |  |              |                                                                                             |
| <b>Parameters</b>  | <p><i>ip-address</i> — Specifies the destination IPv4 or IPv6 address.</p> <table><tr><td><b>Values</b></td><td></td><td></td></tr><tr><td>&lt;ip-address&gt;</td><td>ipv4-address</td><td>a.b.c.d</td></tr><tr><td></td><td>ipv6-address</td><td>x:x:x:x:x:x:x (eight 16-bit pieces)<br/>x:x:x:x:x:d.d.d.d<br/>x - [0..FFFF]H<br/>d - [0..255]D</td></tr></table>                                                                                                                        | <b>Values</b>                                                                               |  |  | <ip-address> | ipv4-address | a.b.c.d |  | ipv6-address | x:x:x:x:x:x:x (eight 16-bit pieces)<br>x:x:x:x:x:d.d.d.d<br>x - [0..FFFF]H<br>d - [0..255]D |
| <b>Values</b>      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                             |  |  |              |              |         |  |              |                                                                                             |
| <ip-address>       | ipv4-address                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | a.b.c.d                                                                                     |  |  |              |              |         |  |              |                                                                                             |
|                    | ipv6-address                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | x:x:x:x:x:x:x (eight 16-bit pieces)<br>x:x:x:x:x:d.d.d.d<br>x - [0..FFFF]H<br>d - [0..255]D |  |  |              |              |         |  |              |                                                                                             |

## dscp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dscp</b> <i>dscp-name</i><br><b>no dscp</b>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>service>if>ies>sap<br>config>service>if>vprn>sap>ip-tunnel                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command sets the DSCP code-point in the outer IP header of encapsulated packets associated with a particular tunnel. The default, set using the no form of the command, is to copy the DSCP value from the inner IP header (after remarking by the private tunnel SAP egress qos policy) to the outer IP header.                                                                                                                                                  |
| <b>Default</b>     | no dscp                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><i>dscp</i> — Specifies the DSCP code-point to be used.</p> <p><b>Values</b></p> <p>be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63</p> |

---

## gre-header

|                    |                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>gre-header</b><br><b>gre-header send-key <i>send-key</i> receive-key <i>receive-key</i></b><br><b>no gre-header</b>                                                                                                                                                                        |
| <b>Context</b>     | config>service>if>ies>sap>ip-tunnel<br>config>service>if>vprn>sap>ip-tunnel                                                                                                                                                                                                                   |
| <b>Description</b> | This command configures the type of the IP tunnel. If the gre-header command is configured then the tunnel is a GRE tunnel with a header inserted between the outer and inner IP headers.<br><br>If the <b>no</b> form of the command is configured then the tunnel is a simple IP-IP tunnel. |
| <b>Default</b>     | no gre-header                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <i>send-key</i> — Specifies the send key value.<br><b>Values</b> 0 to 4294967295<br><i>receive-key</i> — Specifies the receive key value.<br><b>Values</b> 0 to 4294967295                                                                                                                    |

## source

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>source <i>ip-address</i></b><br><b>no source</b>                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>service>if>ies>sap<br>config>service>if>vprn>sap>ip-tunnel                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | This command sets the source IPv4 address of encapsulated packets associated with a particular tunnel. It must be an address in the subnet of the associated public tunnel SAP interface. The GRE does not come up until a valid source address is configured.<br><br>The <b>no</b> form of the command deletes the source address from the tunnel configuration. The tunnel must be administratively shutdown before issuing the <b>no source</b> command. |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies the source IPv4 address of the tunnel.<br><b>Values</b> 1.0.0.0 to 223.255.255.255                                                                                                                                                                                                                                                                                                                                            |

## remote-ip

|               |                                                           |
|---------------|-----------------------------------------------------------|
| <b>Syntax</b> | <b>remote-ip <i>ip-address</i></b><br><b>no remote-ip</b> |
|---------------|-----------------------------------------------------------|



---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>if>ies>sap<br>config>service>if>vprn>sap>ip-tunnel                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | This command sets the primary destination IPv4 address of encapsulated packets associated with a particular tunnel. If this address is reachable in the delivery service (there is a route) then this is the destination IPv4 address of encapsulated packets sent by the delivery service.<br><br>The <b>no</b> form of the command deletes the destination address from the tunnel configuration. |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies the destination IPv4 address of the tunnel.                                                                                                                                                                                                                                                                                                                           |
| <b>Values</b>      | 1.0.0.0 to 223.255.255.255                                                                                                                                                                                                                                                                                                                                                                          |

## frf-12

|                    |                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] frf-12                                                                                       |
| <b>Context</b>     | config>service>ies>if>sap>frame-relay<br>config>service>vprn>if>sap>frame-relay                   |
| <b>Description</b> | This command defines the context to configure the parameters of FRF.12 Frame Relay fragmentation. |

## ete-fragment-threshold

|                    |                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ete-fragment-threshold</b> <i>fragment-threshold</i><br><b>no ete-fragment-threshold</b>                                                                                                                                                      |
| <b>Context</b>     | config>service>ies>if>sap>frame-relay>frf.12<br>config>service>vprn>if>sap>frame-relay>frf.12                                                                                                                                                    |
| <b>Description</b> | This command sets the maximum length, in bytes, of a fragment transmitted across a Frame Relay SAP with the FRF.12 end-to-end fragmentation enabled.<br><br>The no form of this command resets the fragment threshold back to the default value. |
| <b>Default</b>     | ete-fragment-threshold 128                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <i>fragment-threshold</i> — Specifies the maximum fragment length, in bytes, to be transmitted across the FR SAP.                                                                                                                                |
| <b>Values</b>      | 128 to 512 bytes                                                                                                                                                                                                                                 |

## interleave

|               |                                           |
|---------------|-------------------------------------------|
| <b>Syntax</b> | <b>interleave</b><br><b>no interleave</b> |
|---------------|-------------------------------------------|

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>ies>if>sap>frame-relay>frf.12                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command enables interleaving of high priority frames and low-priority frame fragments within a FR SAP using FRF.12 end-to-end fragmentation.</p> <p>When this option is enabled, only frames of the FR SAP non expedited forwarding class queues are subject to fragmentation. The frames of the FR SAP expedited queues are interleaved, with no fragmentation header, among the fragmented frames. In effect, this provides a behavior like in MLPPP Link Fragment Interleaving (LFI).</p> <p>When this option is disabled, frames of all the FR SAP forwarding class queues are subject to fragmentation. The fragmentation header is however not included when the frame size is smaller than the user configured fragmentation size. In this mode, the SAP transmits all fragments of a frame before sending the next full or fragmented frame.</p> <p>The receive direction of the FR SAP supports both modes of operation concurrently, with and without fragment interleaving.</p> <p>The <b>no</b> form of this command restores the default mode of operation.</p> |
| <b>Default</b>     | no interleave                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## scheduling-class

|                    |                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>scheduling-class</b> <i>class-id</i><br><b>no scheduling-class</b>                                                                                                                                                                                     |
| <b>Context</b>     | config>service>ies>if>sap>frame-relay<br>config>service>vprn>if>sap>frame-relay                                                                                                                                                                           |
| <b>Description</b> | This command assigns a Frame Relay scheduling class for a Frame Relay SAP. The scheduling class dictates which queue the frame or frame fragments are stored in FRF.12 end-to-end fragmentation, FRF.12 UNI/NNI link fragmentation and MLFR applications. |
| <b>Default</b>     | scheduling-class 3                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <i>class-id</i> — Specifies the Frame Relay scheduling class number.                                                                                                                                                                                      |
| <b>Values</b>      | 0 to 3                                                                                                                                                                                                                                                    |

## fwd-wholesale

|                    |                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fwd-wholesale</b>                                                                                                                                                                  |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap                                                                                                                                                  |
| <b>Description</b> | This command enables the context to select specific protocols ingressing on the SAP to be redirected to another service. The command is applicable to static SAPs as well as PW-SAPs. |

## fwd-wholesale

|                    |                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] fwd-wholesale</b>                                                                                                                                                                                                                                |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap<br>config>service>vprn>sub-if>grp-if>sap                                                                                                                                                                            |
| <b>Description</b> | This command enables the context to select specific protocols ingressing on the SAP to be redirected to another service. The command is applicable to static SAPs as well as PW-SAPs.<br><br>The <b>no</b> form of this command removes the redirection. |

## pppoe

|                    |                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>pppoe service-id</b><br><b>no pppoe</b>                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>service>ies>if>sap<br>config>service>ies>sub-if>grp-if>sap>fwd-wholesale<br>config>service>vprn>sub-if>grp-if>sap>fwd-wholesale                                                                                                                                                                                           |
| <b>Description</b> | This command specifies that PPPoE packets on ingress on Ethertypes 0x8863 and 0x8864 will be redirected to the specified service. The service referred to by <i>svc-id</i> must be an Epipe service. Redirection to VC-switching Epipe services is not supported.<br><br>The <b>no</b> form of the command removes the redirect. |
| <b>Parameters</b>  | <i>service-id</i> — Specifies the service ID of the Epipe to which packets are redirected.<br><br><b>Values</b> 1 to 2147483647   <i>svc-name</i> up to 64 characters                                                                                                                                                            |

## host-lockout-policy

|                    |                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>host-lockout-policy policy-name</b><br><b>no host-lockout-policy</b>                                                                 |
| <b>Context</b>     | config>service>ies>if>sap                                                                                                               |
| <b>Description</b> | This command configures a host lockout policy.<br><br>The <b>no</b> form of the command removes the policy name from the configuration. |

## host-shutdown

|               |                           |
|---------------|---------------------------|
| <b>Syntax</b> | <b>[no] host-shutdown</b> |
|---------------|---------------------------|

---

|                    |                                                                  |
|--------------------|------------------------------------------------------------------|
| <b>Context</b>     | config>service>ies>if>sap                                        |
| <b>Description</b> | This command administratively enables host creation on this SAP. |

## ip-tunnel

|                    |                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ip-tunnel</b> <i>name</i> [ <b>create</b> ]<br><b>no ip-tunnel</b> <i>name</i>                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>service>ies>if>sap                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command is used to configure an IP-GRE or IP-IP tunnel and associate it with a private tunnel SAP within an IES or VPRN service.</p> <p>The <b>no</b> form of the command deletes the specified IP/GRE or IP-IP tunnel from the configuration. The tunnel must be administratively shutdown before issuing the no ip-tunnel command.</p> |
| <b>Default</b>     | No IP tunnels are defined.                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <i>ip-tunnel-name</i> — Specifies the name of the IP tunnel. Tunnel 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.                                                                                            |

## host

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] host ip</b> <i>ip-address</i> [ <b>mac</b> <i>ieee-address</i> ]] [ <b>subscriber</b> <i>sub-ident-string</i> ] [ <b>sub-profile</b> <i>sub-profile-name</i> ] [ <b>sla-profile</b> <i>sla-profile-name</i> ] [ <b>ancp-string</b> <i>ancp-string</i> ]<br><b>no host</b> {[ <b>ip</b> <i>ip-address</i> ] [ <b>mac</b> <i>ieee-address</i> ]}<br><b>no host all</b>                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>service>ies>if>sap<br>config>service>ies>sub-if>grp-if>sap                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command creates a static subscriber host for the SAP. Static subscriber hosts may be used by the system for various purposes. Applications within the system that make use of static host entries include anti-spoof filters and ARP cache population.</p> <p>Multiple static hosts may be defined on the SAP. Each host is identified by either a source IP address, a source MAC address or both a source IP and source MAC address. Every static host definition must have at least one address defined, IP or MAC.</p> <p>Static hosts can exist on the SAP even with anti-spoof and ARP populate features disabled. When enabled, each feature has different requirements for static hosts.</p> |

**anti-spoof** – When enabled, this feature uses static and dynamic host information to populate entries into an anti-spoof filter table. The anti-spoof filter entries generated will be of the same type as specified in the anti-spoof type parameter. If the SAP anti-spoof filter is defined as **ip**, each static host definition must specify an IP address. If the SAP anti-spoof filter is defined as **ip-mac**, each static host definition must specify both an IP address and MAC address. If definition of a static host is attempted without the appropriate addresses specified for the enabled anti-spoof filter, the static host definition fails.

**arp-populate** – When enabled, this feature uses static and dynamic host information to populate entries in the system ARP cache.

Attempting to define a static subscriber host that conflicts with an existing DHCP Lease State Table entry fails.

Use the **no** form of the command to remove a static entry from the system. The specified *ip-address* and *mac-address* must match the host's exact IP and MAC addresses as defined when it was created. When a static host is removed from the SAP, the corresponding anti-spoof entry and/or ARP cache entry is also removed.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | none                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b> | <p><b>ip</b> <i>ip-address</i> — Specify this optional parameter when defining a static host. The IP address must be specified for <b>anti-spoof ip</b>, <b>anti-spoof ip-mac</b> and <b>arp-populate</b>. Only one static host may be configured on the SAP with a given IP address.</p> <p><b>mac</b> <i>mac-address</i> — Specify this optional parameter when defining a static host. The MAC address must be specified for <b>anti-spoof ip-mac</b> and <b>arp-populate</b>. Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address.</p> <p><b>subscriber</b> <i>sub-ident-string</i> — Specify this optional parameter to specify an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the <b>config&gt;subscr-mgmt&gt;sub-ident-policy</b> context. The subscriber information is used by the SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.</p> <p>For VPRN SAPs with <b>arp-reply-agent</b> enabled with the optional <i>sub-ident</i> parameter, the static subscriber host's sub-ident-string is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPRN destinations.</p> <p>If the static subscriber host's <i>sub-ident</i> string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.</p> <p>If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.</p> <p>If <i>sub-ident</i> is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.</p> <p>ARP requests are never forwarded back to the same SAP or within the receiving SAP's split horizon group.</p> |

**sub-profile** *sub-profile-name* — Specify this optional parameter to specify an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

**sla-profile** *sla-profile-name* — Specify this optional parameter to specify an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

**ancp-string** *ancp-string* — Specifies the ASCII string of the DSLAM circuit ID name.

## 2.5.2.5 Redundant Interface Commands

### redundant-interface

|                    |                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>redundant-interface</b> <i>ip-int-name</i> [ <b>create</b> ]<br><b>no redundant-interface</b> <i>ip-int-name</i>                                                                                                                                                                                                           |
| <b>Context</b>     | config>service>ies                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | This command configures a redundant interface.                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <i>ip-int-name</i> — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.<br><br><b>create</b> — Keyword used to create a redundant interface. |

### address

|                    |                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>address</b> { <i>ip-address/mask</i>   <i>ip-address netmask</i> } [ <b>remote-ip</b> <i>ip-address</i> ]<br><b>no address</b>                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>service>ies>redundant-interface                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | This command assigns an IP address mask or netmask and a remote IP address to the interface.                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <i>ip-address/mask</i> — Assigns an IP address/IP subnet format to the interface.<br><i>ip-address netmask</i> — 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.<br>Assigns an IP address netmask to the interface.<br><b>remote-ip</b> <i>ip-address</i> — Assigns a remote IP to the interface. |

---

## 2.5.2.6 IES Subscriber Interface Commands

### subscriber-interface

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] subscriber-interface</b> <i>ip-int-name</i>                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>service>ies                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command allows the operator to create special subscriber-based interfaces. It is used to contain multiple group interfaces. Multiple subnets associated with the subscriber interface can be applied to any of the contained group interfaces in any combination. The subscriber interface allows subnet sharing between group interfaces.</p> <p>Use the <b>no</b> form of the command to remove the subscriber interface.</p> |
| <b>Parameters</b>  | <i>ip-int-name</i> — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.                                                                                                                                                                                |

### group-interface

|                    |                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>group-interface</b> <i>ip-int-name</i> [create]<br><b>group-interface</b> <i>ip-int-name</i> [create] lns<br><b>group-interface</b> <i>ip-int-name</i> [create] wlangw<br><b>no group-interface</b> <i>ip-int-name</i>                                                                                                              |
| <b>Context</b>     | config>service>ies>subscriber-interface                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command creates a group interface. This interface is designed for triple-play services where multiple SAPs are part of the same subnet. A group interface may contain one or more SAPs.</p> <p>Use the <b>no</b> form of the command to remove the group interface from the subscriber interface.</p>                          |
| <b>Default</b>     | no group interface                                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <i>ip-int-name</i> — Specifies the interface name of a group interface. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.<br><br><b>lns</b> — Specifies that LNS will be used.<br><br><b>wlangw</b> — Specifies that WLAN gateway interface will be used. |

---

## authentication-policy

|                    |                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>authentication-policy</b> <i>name</i><br><b>no authentication-policy</b>                                                                                                            |
| <b>Context</b>     | config>service>ies>if<br>config>service>ies>sub-if>grp-if                                                                                                                              |
| <b>Description</b> | <p>This command assigns an authentication policy to the interface.</p> <p>The <b>no</b> form of this command removes the policy name from the group interface configuration.</p>       |
| <b>Default</b>     | no authentication-policy                                                                                                                                                               |
| <b>Parameters</b>  | <i>name</i> — Specifies the authentication policy name. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. |

## oper-up-while-empty

|                    |                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] oper-up-while-empty</b>                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies>sub-if>grp-if                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command allows the subscriber interface to treat this group interface to be operationally enabled without any active SAPs.</p> <p>This command is typically used with MSAPs where advertising the subnet prior to having a MSAP dynamically created is needed.</p> |

## srrp

|                    |                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] srrp</b> <i>srrp-id</i>                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>service>ies>sub-if>grp-if                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command creates a Subscriber Router Redundancy Protocol (SRRP) instance on a group IP interface. An SRRP instance manages all subscriber subnets within the group interfaces subscriber IP interface or other subscriber IP interfaces that are associated through a wholesale/retail relationship. Only one unique SRRP instance can be configured per group interface.</p> |



The **no** form of the command removes an SRRP instance from a group IP interface. Once removed, the group interface ignores ARP requests for the SRRP gateway IP addresses that may exist on subscriber subnets associated with the group IP interface. Then the group interface stops routing using the redundant IP interface associated with the group IP interface and will stop routing with the SRRP gateway MAC address. Ingress packets destined to the SRRP gateway MAC will also be silently discarded. This is the same behavior as a group IP interface that is disabled (shutdown).

|                   |                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no srrp                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b> | <i>srrp-id</i> — Specifies a 32 bit instance ID that must be unique to the system. The instance ID must also match the instance ID used by the remote router that is participating in the same SRRP context. SRRP is intended to perform a function similar to VRRP where adjacent IP hosts within local subnets use a default gateway to access IP hosts on other subnets. |
| <b>Values</b>     | 1 to 4294967295                                                                                                                                                                                                                                                                                                                                                             |

## bfd-enable

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p>[no] <b>bfd-enable</b> <i>svc-id</i> <b>interface</b> <i>interface-name</i> <b>dst-ip</b> <i>ip-address</i></p> <p>[no] <b>bfd-enable</b> <b>interface</b> <i>interface-name</i> <b>dst-ip</b> <i>ip-address</i> <b>name</b> <i>name</i></p> <p>[no] <b>bfd-enable</b> <b>interface</b> <i>interface-name</i> <b>dst-ip</b> <i>ip-address</i></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>srrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This commands assigns a bi-directional forwarding (BFD) session providing heart-beat mechanism for the given VRRP/SRRP instance. There can be only one BFD session assigned to any given VRRP/SRRP instance, but there can be multiple SRRP/VRRP sessions using the same BFD session. If the interface configured with BFD is using a LAG or a spoke-SDP, the BFD transmit and receive intervals need to be set to a minimum of 300 ms.</p> <p>BFD control the state of the associated interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface. The specified interface may not be configured with BFD; when it is, the virtual router will then initiate the BFD session.</p> <p>The <b>no</b> form of this command removes BFD from the configuration.</p> |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><i>svc-id</i> — Specifies the service ID of the interface running BFD. If no <i>svc-id</i> is specified then it indicates that the interface is a network interface in the Base router instance.</p> <p>This variant of the command is only supported in 'classic' configuration-mode (<b>configure system management-interface configuration-mode classic</b>). The <b>bfd-enable interface interface-name dst-ip ip-address name name</b> variant can be used in all configuration modes.</p> <p><b>Values</b> {<i>id</i>   <i>svc-name</i>}</p>                                                                                                                                                                                                                                                                                                                                                                                                                        |

*id:* 1 to 2147483647  
*svc-name:* Specifies an existing service name up to 64 characters (*svc-name* is an alias for input only. The *svc-name* gets replaced with an id automatically by SR OS in the configuration)

**interface** *interface-name* — Specifies the name of the interface running BFD, up to 32 characters.

**dst-ip** *ip-address* — Specifies the destination address for the BFD session.

**name** *name* — Specifies a service name, up to 64 characters.

## gw-mac

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>gw-mac</b> <i>mac-address</i><br><b>no gw-mac</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>srrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command overrides the default SRRP gateway MAC address used by the SRRP instance. Unless specified, the system uses the same base MAC address for all SRRP instances with the last octet overridden by the lower 8 bits of the SRRP instance ID. The same SRRP gateway MAC address should be in-use by both the local and remote routers participating in the same SRRP context.</p> <p>One reason to change the default SRRP gateway MAC address is if two SRRP instances sharing the same broadcast domain are using the same SRRP gateway MAC. The system will use the SRRP instance ID to separate the SRRP messages (by ignoring the messages that does not match the local instance ID), but a unique SRRP gateway MAC is essential to separate the routed packets for each gateway IP address.</p> <p>The <b>no</b> form of the command removes the explicit SRRP gateway MAC address from the SRRP instance. The SRRP gateway MAC address can only be changed or removed when the SRRP instance is shutdown.</p> |
| <b>Parameters</b>  | <p><i>mac-address</i> — Specifies a MAC address that is used to override the default SRRP base MAC address.</p> <p><b>Values</b> Any MAC address except all zeros, broadcast or multicast addresses. The offset is expressed in normal Ethernet MAC address notation. The defined gw-mac cannot be 00:00:00:00:00:00, ff:ff:ff:ff:ff:ff or any multicast address.</p> <p>If not specified, the system uses the default SRRP gateway MAC address with the last octet set to the 8 least significant bits of the SRRP instance ID.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## keep-alive-interval

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>keep-alive-interval</b> <i>interval</i><br><b>no keep-alive-interval</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>srrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command defines the interval between SRRP advertisement messages sent when operating in the master state. The interval is also the basis for setting the master-down timer used to determine when the master is no longer sending. The system uses three times the keep-alive interval to set the timer. Every time an SRRP advertisement is seen that is better than the local priority, the timer is reset. If the timer expires, the SRRP instance assumes that a master does not exist and initiates the attempt to become master.</p> <p>When in backup state, the SRRP instance takes the keep-alive interval of the master as represented in the masters SRRP advertisement message. Once in master state, the SRRP instance uses its own configured keep-alive interval.</p> <p>The keep-alive-interval may be changed at anytime, but will have no effect until the SRRP instance is in the master state.</p> <p>The <b>no</b> form of the command restores the default interval.</p> |
| <b>Parameters</b>  | <p><i>interval</i> — Specifies the interval, in milliseconds, between SRRP advertisement messages sent when operating in the master state.</p> <p><b>Values</b> 1 to 100</p> <p><b>Default</b> 10 milliseconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## message-path

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>message-path</b> <i>sap-id</i><br><b>no message-path</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>srrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command defines a specific SAP for SRRP in-band messaging. A message-path SAP must be defined prior to activating the SRRP instance. The defined SAP must exist on the SRRP instances group IP interface for the command to succeed and cannot currently be associated with any dynamic or static subscriber hosts. Once a group IP interface SAP has been defined as the transmission path for SRRP Advertisement messages, it cannot be administratively shutdown, will not support static or dynamic subscriber hosts and cannot be removed from the group IP interface.</p> <p>The SRRP instance message-path command may be executed at anytime on the SRRP instance. Changing the message SAP fails if a dynamic or static subscriber host is associated with the new SAP. Once successfully changed, the SRRP instance will immediately disable anti-spoof on the SAP and start sending SRRP Advertisement messages if the SRRP instance is activated.</p> |

Changing the current SRRP message SAP on an active pair of routers should be done in the following manner:

1. Shutdown the backup SRRP instance.
2. Change the message SAP on the shutdown node.
3. Change the message SAP on the active master node.
4. Re-activate the shutdown SRRP instance.

Shutting down the backup SRRP instance prevents the SRRP instances from becoming master due to temporarily using differing message path SAPs.

If an MCS peering is operational between the redundant nodes and the SRRP instance has been associated with the peering, the designated message path SAP will be sent from each member.

The **no** form of the command can only be executed when the SRRP instance is shutdown. Executing no message-path allows the existing SAP to be used for subscriber management functions. A new message-path SAP must be defined prior to activating the SRRP instance.

**Parameters**     *sap-id* — Specifies the physical port identifier portion of the SAP definition.

## policy

**Syntax**     **[no] policy** *vrrp-policy-id*

**Context**     config>service>ies>sub-if>grp-if>srrp

**Description**     This command associates one or more VRRP policies with the SRRP instance. A VRRP policy is a collection of connectivity and verification tests used to manipulate the in-use priorities of VRRP and SRRP instances. A VRRP policy can test the link state of ports, ping IP hosts, discover the existence of routes in the routing table or the ability to reach Layer 2 hosts. When one or more of these tests fail, the VRRP policy has the option of decrementing or setting an explicit value for the in-use priority of an SRRP instance.

More than one VRRP policy may be associated with an SRRP instance. When more than one VRRP policy is associated with an SRRP instance the delta decrement of the in-use priority is cumulative unless one or more test fail that have explicit priority values. When one or more explicit tests fail, the lowest priority value event takes effect for the SRRP instance. When the highest delta-in-use-limit is used to manage the lowest delta derived in-use priority for the SRRP instance.

VRRP policy associations may be added and removed at anytime. A maximum of two VRRP policies can be associated with a single SRRP instance.

The **no** form of the command removes the association with *vrrp-policy-id* from the SRRP instance.

---

**Parameters** *vrrp-policy-id* — Specifies one or more VRRP policies with the SRRP instance.  
**Values** 1 to 9999

## priority

**Syntax** **priority** *priority*  
**no priority**

**Context** config>service>ies>sub-if>grp-if>srrp

**Description** This command overrides the default base priority for the SRRP instance. The SRRP instance priority is advertised by the SRRP instance to its neighbor router and is compared to the priority received from the neighbor router. The router with the best (highest) priority enters the master state while the other router enters the backup state. If the priority of each router is the same, the router with the lowest source IP address in the SRRP advertisement message assumes the master state.

The base priority of an SRRP instance can be managed by VRRP policies. A VRRP policy defines a set of connectivity or verification tests which, when they fail, may lower an SRRP instances base priority (creating an in-use priority for the instance). Every time an SRRP instances in-use priority changes when in master state, it sends an SRRP advertisement message with the new priority. If the dynamic priority drops to zero or receives an SRRP Advertisement message with a better priority, the SRRP instance transitions to the *becoming backup* state. When the priority command is not specified, or the no priority command is executed, the system uses a default base priority of 100. The priority command may be executed at anytime.

The **no** form of the command restores the default base priority to the SRRP instance. If a VRRP policy is associated with the SRRP instance, it will use the default base priority as the basis for any modifications to the SRRP instances in-use priority.

**Parameters** *priority* — Specifies a base priority for the SRRP instance to override the default.  
**Values** 1 to 254

### 2.5.2.7 IES Interface DHCP Commands

## dhcp

**Syntax** **dhcp**

**Context** config>service>ies>if  
config>service>ies>sub-if  
config>service>ies>sub-if>grp-if

---

**Description** This command enters the context to configure DHCP parameters.

## client-applications

**Syntax** **client-applications** {[dhcp] [ppp]}  
**no client-applications**

**Context** config>service>ies>sub-if>dhcp  
config>service>ies>sub-if>grp-if>dhcp

**Description** This command enables the clients that will try to contact the DHCP server(s).  
The **no** form of the command removes the server client type from the configuration.

**Parameters** **dhcp** — Specifies that the DHCP relay will forward requests to the DHCP server(s).  
**ppp** — Specifies that PPPoE will attempt to request an IP address for a PPPoE client from the DHCP server(s).

## action

**Syntax** **action** {replace | drop | keep}  
**no action**

**Context** config>service>ies>if>dhcp>option  
config>service>ies>sub-if>grp-if>dhcp>option

**Description** This command configures the Relay Agent Information Option (Option 82) processing.  
The **no** form of this command returns the system to the default value.

**Default** no action

**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 (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).  
**drop** — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.  
**keep** — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is forwarded towards the client.  
The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.

If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.

## circuit-id

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>circuit-id</b> [ascii-tuple   ifindex   sap-id   vlan-ascii-tuple]<br><b>no circuit-id</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>service>ies>if>dhcp>option<br>config>service>ies>sub-if>grp-if>dhcp>option                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>When enabled, the router sends either an ASCII tuple, or the interface index (If Index), on the specified SAP ID in the <b>circuit-id</b> sub-option of the DHCP packet.</p> <p>If disabled, the <b>circuit-id</b> sub-option of the DHCP packet will be left empty.</p> <p>The <b>no</b> form of this command returns the system to the default.</p>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Default</b>     | circuit-id ascii-tuple                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><b>ascii-tuple</b> — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “ ”.</p> <p><b>ifindex</b> — Specifies that the interface index will be used. The If Index of a router interface can be displayed using the command <b>show&gt;router&gt;if&gt;detail</b>.</p> <p><b>sap-id</b> — Specifies that the SAP ID will be used.</p> <p><b>vlan-ascii-tuple</b> — Specifies that the format will include VLAN ID, dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.</p> |

## match-circuit-id

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] match-circuit-id</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>dhcp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>This command enables Option 82 circuit ID on relayed DHCP packet matching.</p> <p>For Routed CO, the group interface DHCP relay process is stateful. When packets are relayed to the server the virtual router ID, transaction ID, SAP ID, and client hardware MAC address of the relayed packet are tracked. When a response is received from the server the virtual router ID, transaction ID, and client HW MAC address must be matched to determine the SAP on which to send the packet out. In some cases, the virtual router ID, transaction ID, and client hardware MAC address are not guaranteed to be unique.</p> |

When the **match-circuit-id** command is enabled, it is used as part of the key to guarantee correctness in our lookup. This is really only needed when we are dealing with an IP aware DSLAM that proxies the client HW mac address.

**Default** no match-circuit-id

## option

**Syntax** [no] option

**Context** config>service>ies>if>dhcp  
config>service>ies>sub-if>grp-if>dhcp

**Description** This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.

The **no** form of this command returns the system to the default.

**Default** no option

## remote-id

**Syntax** remote-id [mac | string *string*]  
no remote-id

**Context** config>service>ies>if>dhcp>option  
config>service>ies>sub-if>grp-if>dhcp>option

**Description** When enabled, the router sends the MAC address of the remote end (typically the DHCP client) in the **remote-id** sub-option of the DHCP packet. This command identifies the host at the other end of the circuit.

If disabled, the **remote-id** sub-option of the DHCP packet will be left empty.

The **no** form of this command returns the system to the default.

**Default** remote-id

**Parameters** **mac** — Specifies the MAC address of the remote end is encoded in the sub-option.  
**string** *string* — Specifies the remote-id.

## vendor-specific-option

**Syntax** [no] vendor-specific-option

**Context** config>service>ies>if>dhcp>option



---

```
config>service>ies>sub-if>grp-if>dhcp>option
```

**Description** This command configures the vendor specific sub-option of the DHCP relay packet.

## client-mac-address

**Syntax** [no] client-mac-address

**Context** config>service>ies>if>dhcp>option>vendor  
config>service>ies>sub-if>grp-if>dhcp>option>vendor

**Description** This command enables the sending of the MAC address in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address in the vendor specific sub-option of the DHCP relay packet.

## pool-name

**Syntax** [no] pool-name

**Context** config>service>ies>sub-if>grp-if>dhcp>option>vendor

**Description** This command enables the sending of the pool name in the Nokia vendor-specific suboption of the DHCP relay packet.

The **no** form of the command disables the feature.

## sap-id

**Syntax** [no] sap-id

**Context** config>service>ies>if>dhcp>option>vendor  
config>service>ies>sub-if>grp-if>dhcp>option>vendor

**Description** This command enables the sending of the SAP ID in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the SAP ID in the vendor specific sub-option of the DHCP relay packet.

## service-id

**Syntax** [no] service-id

**Context** config>service>ies>if>dhcp>option>vendor

- 
- |                    |                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command enables the sending of the service ID in the vendor specific sub-option of the DHCP relay packet.<br><br>The <b>no</b> form of the command disables the sending of the service ID in the vendor specific sub-option of the DHCP relay packet. |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## string

- |                    |                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>string</b> <i>text</i><br><b>no string</b>                                                                                                                                            |
| <b>Context</b>     | config>service>ies>if>dhcp>option>vendor<br>config>service>ies>sub-if>grp-if>dhcp>option>vendor                                                                                          |
| <b>Description</b> | This command specifies the string in the vendor specific sub-option of the DHCP relay packet.<br><br>The <b>no</b> form of the command returns the default value.                        |
| <b>Parameters</b>  | <i>text</i> — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (" "). |

## system-id

- |                    |                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] system-id</b>                                                                                   |
| <b>Context</b>     | config>service>ies>if>dhcp>option>vendor<br>config>service>ies>sub-if>grp-if>dhcp>option>vendor         |
| <b>Description</b> | This command specifies whether the system-id is encoded in the vendor specific sub-option of Option 82. |

## proxy-server

- |                    |                                                                     |
|--------------------|---------------------------------------------------------------------|
| <b>Syntax</b>      | <b>proxy-server</b>                                                 |
| <b>Context</b>     | config>service>ies>if>dhcp<br>config>service>ies>sub-if>grp-if>dhcp |
| <b>Description</b> | This command configures the DHCP proxy server.                      |

## emulated-server

- |               |                                          |
|---------------|------------------------------------------|
| <b>Syntax</b> | <b>emulated-server</b> <i>ip-address</i> |
|---------------|------------------------------------------|

### no emulated-server

|                    |                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>ies>if>dhcp>proxy-server<br>config>service>ies>sub-if>grp-if>dhcp>proxy-server                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command configures the IP address which will be used as the DHCP server address in the context of this SAP. Typically, the configured address should be in the context of the subnet represented by service.</p> <p>The <b>no</b> form of this command reverts to the default setting. The local proxy server will not become operational without the emulated-server address being specified.</p> |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies the emulated server address.                                                                                                                                                                                                                                                                                                                                                 |

## lease-time

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lease-time</b> [ <i>days days</i> ] [ <i>hrs hours</i> ] [ <i>min minutes</i> ] [ <i>sec seconds</i> ] [ <b>override</b> ]<br><b>no lease-time</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>service>ies>if>dhcp>proxy-server<br>config>service>ies>sub-if>grp-if>dhcp>proxy-server                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command defines the length of lease-time that will be provided to DHCP clients. By default the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.</p> <p>The <b>no</b> form of this command disables the use of the lease-time command. The local-proxy-server will use the lease-time offered by either a RADIUS or DHCP server.</p>                                                                                                                                                                                                                        |
| <b>Default</b>     | 7 days 0 hours 0 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <p><b>override</b> — Specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients.</p> <p><i>days</i> — Specifies the number of days that the given IP address is valid.</p> <p><b>Values</b> 0 to 3650</p> <p><i>hours</i> — Specifies the number of hours that the given IP address is valid.</p> <p><b>Values</b> 0 to 23</p> <p><i>minutes</i> — Specifies the number of minutes that the given IP address is valid.</p> <p><b>Values</b> 0 to 59</p> <p><i>seconds</i> — Specifies the number of seconds that the given IP address is valid.</p> <p><b>Values</b> 0 to 59</p> |

---

## python-policy

|                    |                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>python-policy</b> <i>name</i><br><b>no python-policy</b>                                                                                                        |
| <b>Context</b>     | config>service>ies>if>dhcp<br>config>service>ies>sub-if>dhcp                                                                                                       |
| <b>Description</b> | This command specifies a python policy to be used for DHCPv4. Python policies are configured in the <b>config&gt;python&gt; python-policy</b> <i>name</i> context. |
| <b>Parameters</b>  | <i>name</i> — Specifies the name of an existing python script up to 32 characters in length.                                                                       |

## python-policy

|                    |                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>python-policy</b> <i>name</i><br><b>no python-policy</b>                                                                                                              |
| <b>Context</b>     | config>service>ies>if>dhcp6-relay                                                                                                                                        |
| <b>Description</b> | This command specifies a python policy to be used for DHCPv6 relay. Python policies are configured in the <b>config&gt;python&gt; python-policy</b> <i>name</i> context. |
| <b>Parameters</b>  | <i>name</i> — Specifies the name of an existing python script up to 32 characters in length.                                                                             |

## relay-proxy

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>relay-proxy</b> [ <b>release-update-src-ip</b> ] [ <b>siaddr-override</b> <i>ip-address</i> ]<br><b>no relay-proxy</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>service>ies>if>dhcp<br>config>service>ies>sub-if>dhcp<br>config>service>ies>sub-if>grp-if>dhcp<br>config>service>vprn>if>dhcp<br>config>service>vprn>sub-if>dhcp<br>config>service>vprn>sub-if>grp-if>dhcp                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command enables the DHCPv4 relay proxy function on the interface. The command has no effect when no dhcp servers are configured (DHCPv4 relay not configured). By default, unicast DHCPv4 release messages are forwarded transparently.</p> <p>A relay proxy enhances the relay such that it also relays unicast client DHCPv4 REQUEST messages (lease renewals).</p> <ul style="list-style-type: none"> <li>• In the upstream direction, update the source IP address and add the gateway IP address (gi-address) field before sending the message to the intended DHCP server (the message is not broadcast to all configured DHCP servers).</li> </ul> |

- In the downstream direction, remove the gi-address and update the destination IP address to the address of the `yiaddr` (your IP address) field.

The optional **release-update-src-ip** parameter updates the source IP address of a DHCP RELEASE message with the address used for relayed DHCPv4 messages.

The optional **siaddr-override** *ip-address* parameter enables DHCP server IP address hiding towards the client. This parameter requires that **lease-populate** is enabled on the interface. The DHCP server ip address is required for the address hiding function and is stored in the lease state record. The client interacts with the relay proxy as if it is the DHCP server. In all DHCP messages to the client, the value of following header fields and DHCP options containing the DHCP server IP address is replaced with the configured *<ip-address>*:

- the “source IP address” field in the IP DHCPv4 packet header
- the “siaddr” field in the DHCPv4 header if not equal to zero in the message received from the server
- the Server Identification option (DHCPv4 option 54) if present in the original server message
- the source IP address field in the IP packet header

DHCP OFFER selection during initial binding is done in the relay-proxy. Only the first DHCP OFFER message is forwarded to the client. Subsequent DHCP OFFER messages from different servers are silently dropped.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no relay-proxy                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b> | <p><b>release-update-src-ip</b> — Updates the source IP address of a DHCP RELEASE message with the address used for relayed DHCPv4 messages</p> <p><i>ip-address</i> — Enables DHCPv4 server address hiding towards the DHCPv4 client and activates DHCPv4 OFFER selection in case multiple DHCP servers are configured. The <i>ip-address</i> can be any local address in the same routing instance. If DHCP relay lease-split is enabled, <b>siaddr-override</b> <i>ip-address</i> has priority over the <b>emulated-server</b> <i>ip-address</i> configured in the proxy-server and will be used as the source IP address.</p> |

## server

|                    |                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>server</b> <i>server1</i> [ <i>server2</i> ]                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | <pre>config&gt;service&gt;ies&gt;if&gt;dhcp config&gt;service&gt;ies&gt;sub-if&gt;dhcp config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;dhcp</pre>                                                                                                                                                                                        |
| <b>Description</b> | 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 then the request is forwarded to all of the servers in the list. |

---

|                   |                                                                                                                    |
|-------------------|--------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no server                                                                                                          |
| <b>Parameters</b> | <i>server</i> — Specifies up to eight DHCP server IP address. There can be a maximum of 8 DHCP servers configured. |

## trusted

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] trusted</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>service>ies>if>dhcp<br>config>service>ies>sub-if>grp-if>dhcp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>According to RFC 3046, <i>DHCP Relay Agent Information Option</i>, a DHCP request where the gi-addr is 0.0.0.0 and which contains an Option 82 field in the packet, should be discarded, unless it arrives on a trusted circuit. If trusted mode is enabled on an IP interface, the Relay Agent (the router) will modify the request's gi-addr to be equal to the ingress interface and forward the request.</p> <p>This behavior only applies when the action in the Relay Agent Information Option is "keep". In the case where the Option 82 field is being replaced by the Relay Agent (action = replace), the original Option 82 information is lost anyway, and there is thus no reason for enabling the trusted option.</p> <p>The <b>no</b> form of this command returns the system to the default.</p> |
| <b>Default</b>     | not enabled                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## user-db

|                    |                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>user-db</b> <i>local-user-db-name</i><br><b>no user-db</b>                                                                                                        |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>dhcp                                                                                                                                |
| <b>Description</b> | <p>This command configures the local user database to use for authentication.</p> <p>The <b>no</b> form of the command removes the value from the configuration.</p> |
| <b>Default</b>     | no user-db                                                                                                                                                           |
| <b>Parameters</b>  | <i>local-user-db-name</i> — Specifies the local user database to use for authentication.                                                                             |

## enable-ingress-stats

|                |                                  |
|----------------|----------------------------------|
| <b>Syntax</b>  | <b>[no] enable-ingress-stats</b> |
| <b>Context</b> | config>router>interface          |

```
config>service>ies >interface
config>service>vprn>interface
config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if
```

**Description** This command enables the collection of ingress interface IP stats. This command is only applicable to IP statistics, and not to uRPF statistics.

If enabled, then the following statistics are collected:

- IPv4 offered packets
- IPv4 offered octets
- IPv6 offered packets
- IPv6 offered octets
- Octet statistics for IPv4 and IPv6 bytes at IP interfaces include the layer 2 frame overhead.

**Default** no enable-ingress-stats

## filter

**Syntax** **filter** *filter-id*  
**no filter**

**Context** config>service>ies>sub-if>grp-if>dhcp

**Description** This command configures the DHCP filter for this interface.

**Parameters** *filter-id* — Specifies the filter policy. The filter ID must already exist within the created IP filters.

**Values** 1 to 65535

## gi-address

**Syntax** **gi-address** *ip-address* [*src-ip-addr*]  
**no gi-address**

**Context** config>service>ies>if>dhcp  
config>service>ies>sub-if>grp-if>dhcp

**Description** This command configures the gateway interface address for the DHCP relay. A subscriber interface can include multiple group interfaces with multiple SAPs. The GI address is needed, when the router functions as a DHCP relay, to distinguish between different interfaces.

By default, the GI address used in the relayed DHCP packet is the primary IP address of a normal IES interface. Specifying the GI address allows the user to choose a secondary address. For group interfaces a GI address must be specified under the group interface DHCP context or subscriber-interface DHCP context in order for DHCP to function.

|                   |                                                                                                                                                                                                       |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no gi-address                                                                                                                                                                                         |
| <b>Parameters</b> | <i>ip-address</i> — Specifies the host IP address to be used for DHCP relay packets.<br><i>src-ip-address</i> — Specifies that this GI address is to be the source IP address for DHCP relay packets. |

### 2.5.2.8 PPPoE Commands

#### pppoe

|                    |                                                                |
|--------------------|----------------------------------------------------------------|
| <b>Syntax</b>      | [no] pppoe                                                     |
| <b>Context</b>     | config>service>ies>sub-if>grp-if                               |
| <b>Description</b> | This command enters the context to configure PPPoE parameters. |

#### dhcp-client

|                    |                                                                         |
|--------------------|-------------------------------------------------------------------------|
| <b>Syntax</b>      | dhcp-client                                                             |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>pppoe                                  |
| <b>Description</b> | This command enters the context to configure the PPPoE-to-DHCP options. |

#### ccag-use-origin-sap

|                    |                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] ccag-use-origin-sap                                                                                                                                                                                                                   |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>pppoe>dhcp-client                                                                                                                                                                                         |
| <b>Description</b> | This command enables the original VPLS SAP to be included in the circuit-id option to send to the DHCP server (in case this interface is connected to a VPLS by a CCA MDA).<br><br>The <b>no</b> form of the command disables the feature. |
| <b>Default</b>     | no ccag-use-origin-sap                                                                                                                                                                                                                     |



---

## pap-chap-user-db

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>pap-chap-user-db</b> <i>local-user-db-name</i><br><b>no pap-chap-user-db</b>                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>pppoe                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command configures the local user database to use for PPP Challenge-Handshake Authentication Protocol/Password Authentication Protocol (PAP/CHAP) authentication.</p> <p>If an authentication policy is also configured, <b>pppoe-access-method</b> must be set to none in this authentication policy to use the local user database (in that case RADIUS authentication will not be used for PPPoE hosts).</p> |
| <b>Parameters</b>  | <i>local-user-db-name</i> — Specifies the local user database to use for authentication.                                                                                                                                                                                                                                                                                                                                |

## policy

|                    |                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>policy</b> <i>pppoe-policy-name</i><br><b>no policy</b>                                           |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>pppoe                                                               |
| <b>Description</b> | This command associates a PPPoE policy on this interface.                                            |
| <b>Default</b>     | policy                                                                                               |
| <b>Parameters</b>  | <i>pppoe-policy-name</i> — Specifies a PPPoE policy up to 32 characters in length on this interface. |

## sap-session-limit

|                    |                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sap-session-limit</b> <i>sap-session-limit</i><br><b>no sap-session-limit</b>                            |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>pppoe                                                                      |
| <b>Description</b> | This command specifies the number of PPPoE hosts per SAP allowed for this group-interface.                  |
| <b>Default</b>     | sap-session-limit 1                                                                                         |
| <b>Parameters</b>  | <i>sap-session-limit</i> — Specifies the number of PPPoE hosts per SAP allowed.<br><b>Values</b> 1 to 20000 |

---

## session-limit

|                    |                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>session-limit</b> <i>session-limit</i><br><b>no session-limit</b>                            |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>pppoe                                                          |
| <b>Description</b> | This command specifies the number of PPPoE hosts allowed for this group interface.              |
| <b>Default</b>     | session-limit 1                                                                                 |
| <b>Parameters</b>  | <i>session-limit</i> — Specifies the number of PPPoE hosts allowed.<br><b>Values</b> 1 to 20000 |

## 2.5.2.9 IES Interface ICMP Commands

### hold-time

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>hold-time</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies>interface<br>config>service>ies>subscriber-interface<br>config>service>ies>redundant-interface<br>config>service>vprn>interface<br>config>service>vprn>network-interface<br>config>service>vprn>subscriber-interface<br>config>service>vprn>redundant-interface                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>This command creates the CLI context to configure interface level hold-up and hold-down timers for the associated IP interface.</p> <p>The <b>up</b> timer controls a delay for the associated IPv4 or IPv6 interface so that the system will delay the deactivation of the associated interface for the specified amount of time.</p> <p>The <b>down</b> timer controls a delay for the associated IPv4 or IPv6 interface so that the system will delay the activation of the associated interface for the specified amount of time</p> |

### up

|               |                                                                                                      |
|---------------|------------------------------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>up ip</b> <i>seconds</i><br><b>no up ip</b><br><b>up ipv6</b> <i>seconds</i><br><b>no up ipv6</b> |
|---------------|------------------------------------------------------------------------------------------------------|

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | <pre> config&gt;service&gt;ies&gt;if&gt;hold-time config&gt;service&gt;ies&gt;sub-if&gt;hold-time config&gt;service&gt;ies&gt;red-if&gt;hold-time config&gt;service&gt;vprn&gt;if&gt;hold-time config&gt;service&gt;vprn&gt;nw-if&gt;hold-time config&gt;service&gt;vprn&gt;sub-if&gt;hold-time config&gt;service&gt;vprn&gt;red-if&gt;hold-time </pre>                                                                                                                                                     |
| <b>Description</b> | <p>This command will cause a delay in the deactivation of the associated IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface down.</p> <p>The <b>no</b> form of the command removes the command from the active configuration and removes the delay in deactivating the associated IP interface. If the configuration is removed during a delay period, the currently running delay will continue until it expires.</p> |
| <b>Parameters</b>  | <p><i>seconds</i> — The time delay, in seconds, to make the interface operational.</p> <p><b>Values</b> 1 to 1200</p>                                                                                                                                                                                                                                                                                                                                                                                       |

## down

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <pre> down ip seconds [init-only] no up ip up ipv6 seconds [init-only] no up ipv6 </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | <pre> config&gt;service&gt;ies&gt;if&gt;hold-time config&gt;service&gt;ies&gt;sub-if&gt;hold-time config&gt;service&gt;ies&gt;red-if&gt;hold-time config&gt;service&gt;vprn&gt;if&gt;hold-time config&gt;service&gt;vprn&gt;nw-if&gt;hold-time config&gt;service&gt;vprn&gt;sub-if&gt;hold-time config&gt;service&gt;vprn&gt;red-if&gt;hold-time </pre>                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command will cause a delay in the activation of the associated IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface up, unless the <b>init-only</b> option is configured. If the <b>init-only</b> option is configured, the delay is only applied when the IP interface is first configured or after a system reboot.</p> <p>The <b>no</b> form of the command removes the command from the active configuration and removes the delay in activating the associated IP interface. If the configuration is removed during a delay period, the currently running delay will continue until it completes.</p> |
| <b>Parameters</b>  | <p><i>seconds</i> — The time delay, in seconds, to make the interface operational.</p> <p><b>Values</b> 1 to 1200</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

**init-only** — Specifies that the **down** delay is only applied when the interface is configured or after a reboot.

**Values** 1 to 1200

## icmp

|                    |                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>icmp</b>                                                                                                        |
| <b>Context</b>     | config>service>ies>if<br>config>service>ies>sub-if>grp-if                                                          |
| <b>Description</b> | This command enters the context to configure Internet Control Message Protocol (ICMP) parameters on an IES service |

## mask-reply

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] mask-reply</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>service>ies>if>icmp<br>config>service>ies>sub-if>grp-if>icmp                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.</p> <p>If a local node sends an ICMP mask request to the router interface, the <b>mask-reply</b> command configures the router interface to reply to the request.</p> <p>By default, the router instance will reply to mask requests.</p> <p>The <b>no</b> form of this command disables replies to ICMP mask requests on the router interface.</p> |
| <b>Default</b>     | mask-reply — Reply to ICMP mask requests.                                                                                                                                                                                                                                                                                                                                                                                                                                |

## param-problem

|                    |                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>param-problem</b> [ <i>number seconds</i> ]<br><b>no param-problem</b>                                                                                                                                                                                |
| <b>Context</b>     | config>service>ies>if>icmp<br>config>service>ies>if>icmp6                                                                                                                                                                                                |
| <b>Description</b> | This command specifies whether parameter-problem ICMP messages should be sent. When enabled, parameter-problem ICMP messages are generated by this interface. The <b>no</b> form of the command disables the sending of parameter-problem ICMP messages. |
| <b>Default</b>     | param-problem 100 10                                                                                                                                                                                                                                     |

---

|                   |                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <p><i>number</i> — Specifies the number of parameter-problem ICMP messages to send in the time frame specified by the <i>seconds</i> parameter.</p> <p><b>Values</b> 10 to 1000</p> <p><b>Default</b> 100</p> <p><i>seconds</i> — Specifies the time frame, in seconds, that is used to limit the number of parameter-problem ICMP messages issued.</p> <p><b>Values</b> 1 to 60</p> <p><b>Default</b> 10</p> |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## redirects

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>redirects</b> <i>number seconds</i></p> <p><b>no redirects</b> [<i>number seconds</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | <p>config&gt;service&gt;ies&gt;if&gt;icmp</p> <p>config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;icmp</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.</p> <p>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.</p> <p>The <b>redirects</b> command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional <i>number</i> and <i>seconds</i> parameters by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.</p> <p>The <b>no</b> form of this command disables the generation of icmp redirects on the router interface.</p> |
| <b>Default</b>     | redirects 100 10 — Maximum of 100 redirect messages in 10 seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><i>number</i> — The maximum number of ICMP redirect messages to send. This parameter must be specified with the <i>seconds</i> parameter.</p> <p><b>Values</b> 10 to 1000</p> <p><i>seconds</i> — The time frame in seconds used to limit the <i>number</i> of ICMP redirect messages that can be issued.</p> <p><b>Values</b> 1 to 60</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

---

## ttl-expired

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ttl-expired</b> <i>number seconds</i><br><b>no ttl-expired</b> [ <i>number seconds</i> ]                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>service>ies>if>icmp<br>config>service>ies>sub-if>grp-if>icmp                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.</p> <p>By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.</p> <p>The <b>no</b> form of this command disables the limiting the rate of TTL expired messages on the router interface.</p>                         |
| <b>Default</b>     | ttl-expired 100 10                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <p><i>number</i> — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the <i>seconds</i> parameter.</p> <p><b>Values</b> 10 to 2000</p> <p><i>seconds</i> — The time frame in seconds used to limit the <i>number</i> of ICMP TTL expired messages that can be issued, expressed as a decimal integer.</p> <p><b>Values</b> 1 to 60</p> |

## unreachables

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>unreachables</b> <i>number seconds</i><br><b>no unreachables</b> [ <i>number seconds</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>service>ies>if>icmp<br>config>service>ies>sub-if>grp-if>icmp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command configures the rate for ICMP host and network destination unreachable messages issued on the router interface.</p> <p>The <b>unreachables</b> 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 <i>number</i> and <i>time</i> parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval.</p> <p>By default, generation of ICMP destination unreachable messages is enabled at a maximum rate of 10 per 60 second time interval.</p> <p>The <b>no</b> form of this command disables the generation of icmp destination unreachable messages on the router interface.</p> |

---

|                   |                                                                                                                                              |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | unreachables 100 10                                                                                                                          |
| <b>Parameters</b> | <i>number</i> — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the <i>seconds</i> parameter. |
| <b>Values</b>     | 10 to 2000                                                                                                                                   |
|                   | <i>seconds</i> — The time frame in seconds used to limit the <i>number</i> of ICMP unreachable messages that can be issued.                  |
| <b>Values</b>     | 1 to 60                                                                                                                                      |

## if-attribute

|                    |                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>if-attribute</b>                                                                                                                                         |
| <b>Context</b>     | config>router<br>config>router>interface<br>config>service>ies>interface<br>config>service>vprn>interface                                                   |
| <b>Description</b> | This command creates the context to configure or apply IP interface attributes such as administrative group (admin-group) or Shared Risk Loss Group (SRLG). |

## admin-group

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>admin-group</b> <i>group-name</i> [ <i>group-name</i> ]<br><b>no admin-group</b> <i>group-name</i> [ <i>group-name</i> ]<br><b>no admin-group</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>router>if>if-attribute<br>config>service>ies>if>if-attribute<br>config>service>vprn>if>if-attribute<br>config>router>mpls>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command configures the admin group membership of an interface. The user can apply admin groups to an IES, VPRN, network IP, or MPLS interface. Once an admin group is bound to one or more interface, its value cannot be changed until all bindings are removed.</p> <p>The configured admin-group membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.</p> <p>Only the admin groups bound to an MPLS interface are advertised in TE link TLVs and sub-TLVs when the <b>traffic-engineering</b> option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.</p> |

The **no** form of this command deletes one or more of the admin-group memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.

**Parameters** *group-name* — Specifies up to five groups, each up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain. Each single operation of the **admin-group** command allows a maximum of 5 groups to be specified. However, a maximum of 32 groups can be added to a given interface through multiple operations.

## srlg-group

**Syntax** **srlg-group** *group-name* [*group-name...*(up to 5 max)]  
**no srlg-group** *group-name* [*group-name...*(up to 5 max)]  
**no srlg-group**

**Context** config>router>if>if-attribute  
config>service>ies>if>if-attribute  
config>service>vprn>if>if-attribute  
config>router>mpls>interface

**Description** This command configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface.

An interface can belong to up to 64 SRLG groups. Once an SRLG group is bound to one or more interface, its value cannot be changed until all bindings are removed.

The configured SRLG membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.

Only the SRLGs bound to an MPLS interface are advertised in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

The **no** form of this command deletes one or more of the SRLG memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.

**Parameters** *group-name* — Specifies the name of the group, up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain. Each single operation of the **srlg-group** command allows a maximum of 5 groups to be specified at a time.

## ingress

**Syntax** **ingress**

**Context** config>service>ies>if



**Description** This command enters context to configure ingress parameters for network interfaces.

policy-accounting

**Syntax** **policy-accounting** <template-name>  
**no policy-accounting**

**Context** config>service>ies>if>ingress

**Description** This command configures the service IES interface ingress policy accounting

**Parameters** *template-name* — Specifies the template name, up to 32 characters.

2.5.2.10 IES Interface IPv6 Commands

ipv6

**Syntax** [no] ipv6

**Context** config>service>ies>if

**Description** This command enters the context to configure IPv6 for an IES interface.

address

**Syntax** **address** ipv6-address/prefix-length [eui-64] [track-srrp srrp-instance] [modifier cga-modifier] [dad-disable] [primary-preference primary-preference]  
**no address** ipv6-address/prefix-length

**Context** config>service>ies>if>ipv6

**Description** This command assigns an IPv6 address to the IES interface.

**Parameters** *ipv6-address/prefix-length* — Specifies the IPv6 address on the interface.

| Values               |               |                                                                                             |
|----------------------|---------------|---------------------------------------------------------------------------------------------|
| ipv6-address/prefix: | ipv6-address  | x:x:x:x:x:x:x (eight 16-bit pieces)<br>x:x:x:x:x:d.d.d.d<br>x [0 to FFFF]H<br>d [0 to 255]D |
|                      | prefix-length | 1 to 128                                                                                    |

**eui-64** — When the **eui-64** keyword is specified, a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from MAC address on Ethernet interfaces. For interfaces without a MAC address, for example ATM interfaces, the Base MAC address of the chassis is used.

**track-srrp** *srrp-instance* — Specifies the SRRP instance ID that this interface route needs to track.

**Values** 1 to 4294967295

**modifier** *cga-modifier* — Specifies the modifier in 32 hexadecimal nibbles.

**Values** 0x0–0xFFFFFFFF

**dad-disable** — Disables Duplicate Address Detection (DAD) and sets the address to preferred, even if there is a duplicated address.

**primary-preference** *primary-preference* — Specifies a *primary-preference* index to an IPv6 address of the interface to enforce the order in which the address is used by control plane protocols and applications which require a fixed address of the interface. These include LDP and Segment Routing.

When originating packets from this interface, the source IPv6 address follows the selection rules in RFC 6724 except for the specific cases where a fixed address is required. In the latter case, the IPv6 address with the lowest primary-preference index is selected. If the selected address is removed, the system selects the IPv6 address with the next lowest *primary-preference* index.

The system assigns the next available index value to any IPv6 address of the interface when configured without the *primary-preference* index value specified. The address index space is unique across all addresses of a given interface.

**Values** 1 to 4294967295

## dad-disable

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] dad-disable</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>if>ipv6                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command disables duplicate address detection (DAD) on a per-interface basis. This prevents the router from performing a DAD check on the interface. All IPv6 addresses of an interface with DAD disabled, immediately enter a preferred state, without checking for uniqueness on the interface. This is useful for interfaces which enter a looped state during troubleshooting and operationally disable themselves when the loop is detected, requiring manual intervention to clear the DAD violation.</p> <p>The <b>no</b> form of the command turns off <b>dad-disable</b> on the interface.</p> |
| <b>Default</b>     | no dad-disable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## dhcp6-relay

|                    |                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] dhcp6-relay</b>                                                                                                                                     |
| <b>Context</b>     | config>service>ies>if>ipv6                                                                                                                                  |
| <b>Description</b> | This command enters the context to configure DHCPv6 relay parameters for the IES interface.<br><br>The <b>no</b> form of the command disables DHCPv6 relay. |

## lease-populate

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lease-populate</b> <i>[nbr-of-leases]</i><br><b>lease-populate</b> <i>[nbr-of-leases]</i> <b>route-populate</b> <b>[pd]</b> <b>na</b> <b>[ta]</b><br><b>lease-populate</b> <i>[nbr-of-leases]</i> <b>route-populate</b> <b>pd</b> <b>[na]</b> <b>[ta]</b> <b>[exclude]</b><br><b>lease-populate</b> <i>[nbr-of-leases]</i> <b>route-populate</b> <b>[pd]</b> <b>[na]</b> <b>ta</b><br><b>no lease-populate</b>                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>service>ies>if>ipv6>dhcp-relay                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command specifies the maximum number of DHCPv6 lease states allocated by the DHCPv6 relay function, allowed on this interface.</p> <p>Optionally, by specifying “route-populate” parameter, system could:</p> <ul style="list-style-type: none"> <li>• Create routes based on the IA_PD/IA_NA/IA_TA prefix option in relay-reply message.</li> <li>• Create black hole routes based on OPTION_PD_EXCLUDE in IA_PD in relay-reply message.</li> </ul> <p>These routes could be redistributed into IGP/BGP by using route-policy, following protocol types that could be used in “from protocol”:</p> <ul style="list-style-type: none"> <li>• dhcpv6-pd</li> <li>• dhcpv6-na</li> <li>• dhcpv6-ta</li> <li>• dhcpv6-pd-excl</li> </ul> |
| <b>Parameters</b>  | <p><i>nbr-of-entries</i> — Defines the number lease state table entries allowed for this interface. If this parameter is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCPv6 ACK messages are discarded.</p> <p><b>Values</b> 1 to 8000</p> <p><b>route-populate</b> — Specifies the route populate parameter.</p> <p><b>Values</b> <b>pd/na/ta</b> — Create route based on specified option.<br/> <b>exclude</b> — Create blackhole route based on OPTION_PD_EXCLUDE.</p>                                                                                                                                                     |

---

## link-address

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>link-address</b> <i>ipv6-address</i><br><b>no link-address</b>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>service>ies>if>ipv6>dhcp6-relay                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command configures the link address used for prefix selection at the DHCP server.</p> <p>The link-address is a field in DHCP6 Relay-Forward message that is used in DHCP6 server to select the IPv6 address (IA-NA) or IPv6 prefix (IA-PD) from a pool with configured prefix range covering the link-address. The selection scope is the pool or a prefix range within the pool.</p> <p>The <b>no</b> form of the command reverts to the default.</p> |
| <b>Default</b>     | no link-address                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <i>ipv6-address</i> — Specifies the link-address.                                                                                                                                                                                                                                                                                                                                                                                                              |
|                    | <b>Values</b><br><ipv6-address> <i>ipv6-address</i> - x:x:x:x:x:x:x (eight 16-bit pieces)<br>x:x:x:x:x:d.d.d.d<br>x - [0 to FFFF]H<br>d - [0 to 255]D                                                                                                                                                                                                                                                                                                          |

## neighbor-resolution

|                    |                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] neighbor-resolution</b>                                                                                                           |
| <b>Context</b>     | config>service>ies>if>ipv6>dhcp6-relay                                                                                                    |
| <b>Description</b> | <p>This command enables neighbor resolution with DHCPv6 relay.</p> <p>The <b>no</b> form of the command disables neighbor resolution.</p> |

## option

|                    |                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] option</b>                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>if>ipv6>dhcp6-relay                                                                                                                                    |
| <b>Description</b> | <p>This command enters the context to configure DHCPv6 relay information options.</p> <p>The <b>no</b> form of the command disables DHCPv6 relay information options.</p> |

## interface-id

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>interface-id [ascii-tuple]</b><br><b>interface-id ifindex</b><br><b>interface-id sap-id</b><br><b>interface-id string</b> <i>string</i><br><b>no interface-id</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>service>ies>if>ipv6>dhcp6>option                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command enables the sending of interface ID options in the DHCPv6 relay packet.</p> <p>The <b>no</b> form of the command disables the sending of interface ID options in the DHCPv6 relay packet</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <p><b>ascii-tuple</b> — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “ ”.</p> <p><b>ifindex</b> — Specifies that the interface index will be used (the If Index of a router interface can be displayed using the command <b>show&gt;router&gt;if&gt;detail</b>).</p> <p><b>sap-id</b> — Specifies that the SAP identifier will be used.</p> <p><b>string</b> <i>string</i> — Specifies a string of up to 32 characters long, 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.</p> |

## remote-id

|                    |                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] remote-id</b>                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>service>ies>if>ipv6>dhcp6>option                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command enables the sending of remote ID option in the DHCPv6 relay packet.</p> <p>The client DHCP Unique Identifier (DUID) is used as the remote ID.</p> <p>The <b>no</b> form of the command disables the sending of remote ID option in the DHCPv6 relay packet.</p> |

## server

|                |                                                           |
|----------------|-----------------------------------------------------------|
| <b>Syntax</b>  | <b>server</b> <i>ipv6-address</i> [ <i>ipv6-address</i> ] |
| <b>Context</b> | config>service>ies>if>ipv6>dhcp6                          |

---

|                    |                                                                                                                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command specifies a list of servers where DHCPv6 requests will be forwarded. The list of servers can entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCPv6 relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list. |
| <b>Default</b>     | no server                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <i>ipv6-address</i> — Specifies the IPv6 addresses of the DHCP servers where the DHCPv6 requests will be forwarded. Up to 8 addresses can be specified.                                                                                                                                                                                        |
| <b>Values</b>      | <div> <div>ipv6-address:</div> <div> x:x:x:x:x:x:x (eight 16-bit pieces)<br/> x:x:x:x:x:d.d.d.d<br/> x - [0..FFFF]H<br/> d - [0..255]D </div> </div>                                                                                                                                                                                           |

## source-address

|                    |                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>source-address</b> <i>ipv6-address</i><br><b>no source-address</b>                                                                                    |
| <b>Context</b>     | config>service>ies>if>ipv6>dhcp6                                                                                                                         |
| <b>Description</b> | This command configures the source IPv6 address of the DHCPv6 relay messages.                                                                            |
| <b>Parameters</b>  | <i>ipv6-address</i> — Specifies the source IPv6 address of the DHCPv6 relay messages.                                                                    |
| <b>Values</b>      | <div> <div>ipv6-address:</div> <div> x:x:x:x:x:x:x (eight 16-bit pieces)<br/> x:x:x:x:x:d.d.d.d<br/> x - [0 to FFFF]H<br/> d - [0 to 255]D </div> </div> |

## dhcp6-server

|                    |                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] dhcp6-server</b>                                                                                                                                                 |
| <b>Context</b>     | config>service>ies>if>ipv6                                                                                                                                               |
| <b>Description</b> | <p>This command enters the context to configure DHCPv6 server parameters for the IES interface.</p> <p>The <b>no</b> form of the command disables the DHCPv6 server.</p> |

## max-nbr-of-leases

|                    |                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>max-nbr-of-leases</b> <i>max-nbr-of-leases</i><br><b>no max-nbr-of-leases</b>                                                                                                                                |
| <b>Context</b>     | config>service>ies>if>ipv6>dhcp6-server                                                                                                                                                                         |
| <b>Description</b> | <p>This command configures the maximum number of lease states installed by the DHCPv6 server function allowed on this interface.</p> <p>The <b>no</b> form of the command returns the value to the default.</p> |
| <b>Default</b>     | max-nbr-of-leases 8000                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><i>max-nbr-of-leases</i> — Specifies the maximum number of lease states installed by the DHCPv6 server function allowed on this interface.</p> <p><b>Values</b>      0 to 8000</p>                           |

## prefix-delegation

|                    |                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] prefix-delegation</b>                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>service>ies>if>ipv6>dhcp6-server                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command configures prefix delegation options for delegating a long-lived prefix from a delegating router to a requesting router, where the delegating router does not require knowledge about the topology of the links in the network to which the prefixes will be assigned.</p> <p>The <b>no</b> form of the command disables prefix-delegation.</p> |

prefix

|                      |                                                                                                                 |
|----------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>        | [no] <b>prefix</b> <i>ipv6-address/prefix-length</i>                                                            |
| <b>Context</b>       | config>service>ies>if>ipv6>dhcp6-server>pfx-delegate                                                            |
| <b>Description</b>   | This command specifies the IPv6 prefix that will be delegated by this system.                                   |
| <b>Parameters</b>    | <i>ipv6-address/prefix-length</i> — Specifies the IPv6 address on the interface.                                |
| <b>Values</b>        |                                                                                                                 |
| ipv6-address/prefix: | ipv6-address        x:x:x:x:x:x:x (eight 16-bit pieces)<br>x:x:x:x:x:d.d.d.d<br>x [0 to FFFF]H<br>d [0 to 255]D |
| prefix-length        | 1 to 128                                                                                                        |

---

## duid

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>duid</b> <i>duid</i> [ <b>iaid</b> <i>iaid</i> ]<br><b>no duid</b>                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | This command configures the DHCP Unique Identifier (DUID) of the DHCP client.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><i>duid</i> — Specifies the ID of the requesting router. If set to a non-zero value the prefix defined will only be delegated to this router. If set to zero, the prefix will be delegated to any requesting router.</p> <p><b>iaid</b> <i>iaid</i> — Specifies the identity association identification (IAID) from the requesting router that needs to match in order to delegate the prefix defined in this row. If set to 0 no match on the received IAID is done.</p> |

## preferred-lifetime

|                    |                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>preferred-lifetime</b> <i>seconds</i><br><b>preferred-lifetime</b> <b>infinite</b><br><b>no preferred-lifetime</b>                                                                                                        |
| <b>Context</b>     | config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix                                                                                                                                                                         |
| <b>Description</b> | <p>This command configures the IPv6 prefix/mask preferred life time. The preferred-lifetime value cannot be bigger than the valid-lifetime value.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | 604800 seconds (7 days)                                                                                                                                                                                                      |
| <b>Parameters</b>  | <p><i>seconds</i> — Specifies the time, in seconds, that this prefix remains preferred.</p> <p><b>Values</b>      1 to 4294967294</p> <p><b>infinite</b> — Specifies that this prefix remains preferred infinitely.</p>      |

## valid-lifetime

|                    |                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>valid-lifetime</b> <i>seconds</i><br><b>valid-lifetime</b> <b>infinite</b><br><b>no valid-lifetime</b>  |
| <b>Context</b>     | config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix                                                       |
| <b>Description</b> | This command configures the time, in seconds, that the prefix is valid. 4,294,967,295 represents infinity. |



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

|                   |                                                                                  |
|-------------------|----------------------------------------------------------------------------------|
| <b>Default</b>    | 2592000 seconds (30 days)                                                        |
| <b>Parameters</b> | <i>seconds</i> — Specifies the time, in seconds, that this prefix remains valid. |
| <b>Values</b>     | 1 to 4294967295                                                                  |
|                   | <b>infinite</b> — Specifies that this prefix remains valid infinitely.           |

## icmp6

|                    |                                                                  |
|--------------------|------------------------------------------------------------------|
| <b>Syntax</b>      | <b>icmp6</b>                                                     |
| <b>Context</b>     | config>service>ies>if>ipv6                                       |
| <b>Description</b> | This command configures ICMPv6 parameters for the IES interface. |

## packet-too-big

|                    |                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>packet-too-big</b> [ <i>number seconds</i> ]<br><b>no packet-too-big</b>                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>service>ies>if>ipv6>icmp6                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | This command specifies whether “packet-too-big” ICMPv6 messages should be sent. When enabled, ICMPv6 “packet-too-big” messages are generated by this interface.<br><br>The <b>no</b> form of the command disables the sending of ICMPv6 “packet-too-big” messages.                                                                                                                             |
| <b>Default</b>     | packet-too-big 100 10                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <i>number</i> — Specifies the number of “packet-too-big” ICMPv6 messages to send in the time frame specified by the <i>seconds</i> parameter.<br><br><b>Values</b> 10 to 1000<br><b>Default</b> 100<br><br><i>seconds</i> — Specifies the time frame in seconds that is used to limit the number of “packet-too-big” ICMPv6 messages issued.<br><br><b>Values</b> 1 to 60<br><b>Default</b> 10 |

## redirects

|               |                                                                   |
|---------------|-------------------------------------------------------------------|
| <b>Syntax</b> | <b>redirects</b> [ <i>number seconds</i> ]<br><b>no redirects</b> |
|---------------|-------------------------------------------------------------------|

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>ies>if>ipv6>icmp6                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command configures ICMPv6 redirect messages. When enabled, ICMPv6 redirects are generated when routes are not optimal on this router and another router on the same subnetwork has a better route in order to alert that node that a better route is available.</p> <p>When disabled, ICMPv6 redirects are not generated.</p>                                                        |
| <b>Default</b>     | redirects 100 10                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <p><i>number</i> — Specifies the number of version 6 redirects are to be issued in the time frame specified by the <i>seconds</i> parameter.</p> <p><b>Values</b> 10 to 1000</p> <p><b>Default</b> 100</p> <p><i>seconds</i> — Specifies the time frame in seconds that is used to limit the number of version 6 redirects issued.</p> <p><b>Values</b> 1 to 60</p> <p><b>Default</b> 10</p> |

## time-exceeded

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>time-exceeded</b> [ <i>number seconds</i> ]<br><b>no time-exceeded</b>                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>service>ies>if>ipv6>icmp6                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>This command specifies whether “time-exceeded” ICMPv6 messages should be sent. When enabled, ICMPv6 “time-exceeded” messages are generated by this interface.</p> <p>When disabled, ICMPv6 “time-exceeded” messages are not sent.</p>                                                                                                                                                                                  |
| <b>Default</b>     | time-exceeded 100 10                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <p><i>number</i> — Specifies the number of “time-exceeded” ICMPv6 messages are to be issued in the time frame specified by the <i>seconds</i> parameter.</p> <p><b>Values</b> 10 to 2000</p> <p><b>Default</b> 100</p> <p><i>seconds</i> — Specifies the time frame in seconds that is used to limit the number of “time-exceeded” ICMPv6 message to be issued.</p> <p><b>Values</b> 1 to 60</p> <p><b>Default</b> 10</p> |

## unreachables

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>unreachables</b> [ <i>number seconds</i> ]<br><b>no unreachable</b>                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>service>ies>if>ipv6>icmp6                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command specifies that ICMPv6 host and network unreachable messages are generated by this interface.<br><br>When disabled, ICMPv6 host and network unreachable messages are not sent.                                                                                                                                                                                                                            |
| <b>Default</b>     | unreachables 100 10                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <i>number</i> — Specifies the number of destination unreachable ICMPv6 messages are issued in the time frame specified by the <i>seconds</i> parameter.<br><br><b>Values</b> 10 to 2000<br><b>Default</b> 100<br><br><i>seconds</i> — Specifies the time frame in seconds that is used to limit the number of destination unreachable ICMPv6 messages to be issued.<br><br><b>Values</b> 1 to 60<br><b>Default</b> 10 |

## link-local-address

|                    |                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>link-local-address</b> <i>ipv6-address</i> [ <b>dad-disable</b> ]<br><b>no link-local-address</b>                                                                                                   |
| <b>Context</b>     | config>router>if>ipv6<br>config>service>ies>if>ipv6<br>config>service>vprn>if>ipv6                                                                                                                     |
| <b>Description</b> | This command configures the IPv6 link local address.<br><br>The no form of the command removes the configured link local address, and the router automatically generates a default link local address. |



**Caution:** Removing a manually configured link local address may impact routing protocols or static routes that have a dependency on that address. It is not recommended to remove a link local address when there are active IPv6 subscriber hosts on an IES or VPRN interface.

|                   |                                                                                                                                           |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <b>dad-disable</b> — Disables Duplicate Address Detection (DAD) and sets the address to preferred, even if there is a duplicated address. |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------|

---

## local-proxy-nd

|                    |                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] local-proxy-nd</b>                                                                                                                              |
| <b>Context</b>     | config>service>ies>if>ipv6                                                                                                                              |
| <b>Description</b> | This command enables local proxy neighbor discovery on the interface.<br><br>The <b>no</b> form of the command disables local proxy neighbor discovery. |

## nd-learn-unsolicited

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>nd-learn-unsolicited {global   link-local   both}</b><br><b>no nd-learn-unsolicited</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>service>ies>if>ipv6                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | This command enables the ability to learn neighbor entries out of received unsolicited Neighbor Advertisement messages with or without the solicited flag set. The command can be enabled for global addresses, link-local addresses, or for both.<br><br>The <b>no</b> form of the command makes the router use standard RFC 4861 behavior, as described below, for learning of neighbor entries. <ul style="list-style-type: none"><li>• If an unsolicited NA, regardless of the S flag, is received from a neighbor that is not yet in the ND cache, the NA is ignored.</li><li>• If an NS, RS, RA, or Redirect message with a Link Layer Address (MAC) is received from a neighbor that is not yet in the ND cache, a new neighbor entry is created in the cache to store the received Link Layer MAC. The neighbor is put in the stale state.</li></ul> |
| <b>Parameters</b>  | <b>global</b> — Learns global neighbor entries out of received unsolicited Neighbor Advertisement messages.<br><br><b>link-local</b> — Learns link local neighbor entries out of received unsolicited Neighbor Advertisement messages.<br><br><b>both</b> — Learns both global and link local neighbor entries out of received unsolicited Neighbor Advertisement messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## nd-populate-host-route

|                |                                    |
|----------------|------------------------------------|
| <b>Syntax</b>  | <b>[no] nd-populate-host-route</b> |
| <b>Context</b> | config>service>ies>interface>ipv6  |

**Description** This command enables the addition or deletion of host routes in the route-table derived from neighbor entries in the neighbor cache. To enable this command, the interface must be shut down. The command triggers the population of host routes in the route table out of their corresponding static, dynamic, or EVPN types in the neighbor table. Neighbor entries installed by subscriber management, local interfaces, and others, do not create host-routes.

Only reachable entries are added to the route table (entries are created from solicited NA messages). Entries created as stale — from Neighbor Solicitation (NS), unsolicited Neighbor Advertisements (NA), Router Solicitation (RS), Router Advertisement (RA), and Redirect messages — are not added to the route table because the neighbor is not confirmed as two-way.

- RA, RS, NS, and Redirect messages with a link layer address are added as STALE cache entries. Unsolicited NAs are added as STALE if **nd-learn-unsolicited** is configured.
- To speed up the addition of host routes to the route table for neighbors created as STALE, the following procedure is used:
  - If **nd-populate-host-route** is configured, the router sends an NS (unicast Neighbor Unreachability Detection (NUD) message) to the neighbor created as STALE. Only one NUD message is sent.
  - If **nd-populate-host-route** is not configured, no confirmation message is sent and regular procedures apply.
- When the solicited NA for the neighbor is received, the entry becomes reachable and is then added to the route-table.

The **no** form of this command disables the creation of host routes from the neighbor cache.

## nd-proactive-refresh

**Syntax** **nd-proactive-refresh** {**global** | **link-local** | **both**}  
**no nd-proactive-refresh**

**Context** config>service>ies>if>ipv6

**Description** This command enables a proactive refresh of the neighbor entries. When enabled, at the stale timer expiration, the router sends a NUD message to the host (regardless of the existence of traffic to the IP address on the IOM), so the entry can be refreshed or removed.

This behavior is different from ARP, where the refresh is sent 30 seconds prior to the entry's age out time. The refresh can be optionally enabled for global addresses, link-local addresses, or both.

The **no** form of the command disables the proactive behavior and the router only refreshes an entry if there is traffic that needs to be sent to the IP address.

**Parameters** **global** — Refreshes global neighbor entries.  
**link-local** — Refreshes link local neighbor entries.

---

**both** — Refreshes both global and link local neighbor entries.

## nd-route-tag

|                    |                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>nd-route-tag</b> <i>tag</i><br><b>no nd-route-tag</b>                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>service>ies>if>ipv6                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command adds a route tag to the ARP-ND host routes generated out of the neighbor entries in the interface. As any other route tag, it can be used to match ARP-ND routes in BGP export policies.</p> <p>The <b>no</b> form of this command removes the route tag for the ARP-ND host routes.</p> |
| <b>Parameters</b>  | <p><i>tag</i> — Specifies the route tag to be added when the proxy ND entries are advertised to EVPN.</p> <p><b>Values</b>      1 to 255</p>                                                                                                                                                             |

## proxy-nd-policy

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>proxy-nd-policy</b> <i>policy-name</i> [ <i>policy-name</i> ]<br><b>no proxy-nd-policy</b>                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>service>ies>if>ipv6                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | This command applies a proxy neighbor discovery policy for the interface.                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <p><i>policy-name</i> — Specifies an existing neighbor discovery policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. The specified policy name(s) must already be defined. Up to 5 policy-names can be specified in a single statement.</p> |

## neighbor

|                    |                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>neighbor</b> <i>ipv6-address mac-address</i><br><b>no neighbor</b> <i>ipv6-address</i> |
| <b>Context</b>     | config>service>ies>if>ipv6                                                                |
| <b>Description</b> | This command configures IPv6-to-MAC address mapping on the IES interface.                 |
| <b>Default</b>     | n/a                                                                                       |

**Parameters** *ipv6-address* — The IPv6 address of the interface for which to display information.

**Values**

*ipv6-address*: x:x:x:x:x:x:x (eight 16-bit pieces)  
x:x:x:x:x:x.d.d.d  
x - [0..FFFF]H  
d - [0..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 non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

## neighbor-limit

**Syntax** **neighbor-limit** *limit* [**log-only**] [**threshold percent**]

**no neighbor-limit**

**Context** config>service>ies>if>ipv6

**Description** This command configures the maximum amount of dynamic IPv6 neighbor entries that can be learned on an IP interface.

When the number of dynamic neighbor entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations will be dropped. Entries that have already been learned will be refreshed.

The **no** form of the command removes the **neighbor-limit**.

**Default** no neighbor-limit

**Parameters** **log-only** — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit will be learned.

*percent* — The threshold value (as a percentage) that triggers a warning message to be sent.

**Values** 0 to 100

*limit* — The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic neighbor learning is disabled and no dynamic neighbor entries are learned.

**Values** 0 to 102400

---

## backup

|                    |                                                                        |
|--------------------|------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] backup</b> <i>ip-address</i>                                   |
| <b>Context</b>     | config>service>ies>if>ipv6>vrrp                                        |
| <b>Description</b> | This command configures virtual router IP addresses for the interface. |

## init-delay

|                    |                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>init-delay</b> <i>seconds</i><br><b>no init-delay</b>                                                    |
| <b>Context</b>     | config>service>ies>if>ipv6>vrrp                                                                             |
| <b>Description</b> | This command configures a VRRP initialization delay timer.                                                  |
| <b>Default</b>     | no init-delay                                                                                               |
| <b>Parameters</b>  | <i>seconds</i> — Specifies the initialization delay timer for VRRP, in seconds.<br><b>Values</b> 1 to 65535 |

## mac

|                    |                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mac</b> <i>mac-address</i><br><b>no mac</b>                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>service>ies>if>ipv6>vrrp                                                                                                                                                                                                                                                    |
| <b>Description</b> | This command assigns a specific MAC address to an IES IP interface.<br><br>The <b>no</b> form of the command returns the MAC address of the IP interface to the default value.                                                                                                     |
| <b>Default</b>     | The physical MAC address associated with the Ethernet interface that the SAP is configured on (the default MAC address assigned to the interface, assigned by the system).                                                                                                         |
| <b>Parameters</b>  | <i>mac-address</i> — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses. |

## master-int-inherit

|                |                                 |
|----------------|---------------------------------|
| <b>Syntax</b>  | <b>[no] master-int-inherit</b>  |
| <b>Context</b> | config>service>ies>if>ipv6>vrrp |



**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>ipv6>vrrp

**Description** This command sets the advertisement timer and indirectly sets the master down timer on the virtual router instance. The message-interval setting must be the same for all virtual routers participating as a virtual router. Any VRRP advertisement message received with an Advertisement Interval field different than the virtual router instance configured message-interval value will be silently discarded.

The message-interval command is available in both non-owner and owner **vrrp** *virtual-router-id* nodal contexts. If the message-interval command is not executed, the default message interval of 1 second will be used.

The **no** form of this command restores the default message interval value of 1 second to the virtual router instance.

**Parameters** *seconds* — The number of seconds that will transpire before the advertisement timer expires.

**Values** 1 to 255

**Default** 1

**milliseconds** *milliseconds* — Specifies the time interval, in milliseconds, between sending advertisement messages.

**Values** 100 to 900

## oper-group

**Syntax** **oper-group** *group-name*  
**no oper-group**

**Context** config>service>ies>if>vrrp

**Description** This command configures VRRP to associate with an operational group. When associated, VRRP notifies the operational group of its state changes so that other protocols can monitor it to provide a redundancy mechanism. When VRRP is the master router (MR), the operational group is up and is down for all other VRRP states.

The **no** form of the command removes the association.

---

|                   |                                                                                               |
|-------------------|-----------------------------------------------------------------------------------------------|
| <b>Default</b>    | no oper-group — No operational group is configured.                                           |
| <b>Parameters</b> | <i>group-name</i> — Specifies the operational group identifier up to 32 characters in length. |

## ping-reply

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ping-reply</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command enables the non-owner master to reply to ICMP echo requests directed at the virtual router instances IP addresses. The ping request can be received on any routed interface.</p> <p>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.</p> <p>Non-owner backup virtual routers never respond to ICMP echo requests regardless of the setting of ping-reply configuration.</p> <p>The ping-reply command is only available in non-owner <b>vrrp</b> <i>virtual-router-id</i> nodal context. If the ping-reply command is not executed, ICMP echo requests to the virtual router instance IP addresses will be silently discarded.</p> <p>The <b>no</b> form of this command restores the default operation of discarding all ICMP echo request messages destined to the non-owner virtual router instance IP addresses.</p> |
| <b>Default</b>     | no ping-reply                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## policy

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>policy vrrp-policy-id</b><br><b>no policy</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>service>ies>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command creates VRRP control policies. The VRRP policy ID must be created by the policy command prior to association with the virtual router instance.</p> <p>The policy command provides the ability to associate a VRRP priority control policy to a virtual router instance. The policy may be associated with more than one virtual router instance. The priority events within the policy either override or diminish the base-priority dynamically affecting the in-use priority. As priority events clear in the policy, the in-use priority may eventually be restored to the base-priority value.</p> |

The policy command is only available in the non-owner **vrrp** *virtual-router-id* nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed by VRRP priority control policies. For non-owner virtual router instances, if the policy command is not executed, the base-priority will be used as the in-use priority.

The **no** form of this command removes any existing VRRP priority control policy association from the virtual router instance. All such associations must be removed prior to the policy being deleted from the system.

|                   |                                                                                                                                                                                                                                                  |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | n/a                                                                                                                                                                                                                                              |
| <b>Parameters</b> | <i>vrrp-policy-id</i> — The vrrp-policy-id parameter associated the corresponding VRRP priority control policy-id with the virtual router instance. The vrrp-policy-id must already exist in the system for the policy command to be successful. |
| <b>Values</b>     | 1 to 9999                                                                                                                                                                                                                                        |

## preempt

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] preempt</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>service>ies>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>The preempt mode value controls whether a specific backup virtual router preempts a lower priority master.</p> <p>When preempt is enabled, the virtual router instance overrides any non-owner master with an “in use” message priority value less than the virtual router instance in-use priority value. If preempt is disabled, the virtual router only becomes master if the master down timer expires before a VRRP advertisement message is received from another virtual router.</p> <p>The IP address owner will always become master when available. Preempt mode cannot be disabled on the owner virtual router.</p> <p>The default value for preempt mode is enabled.</p> |
| <b>Default</b>     | preempt                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## priority

|                    |                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>priority</b> <i>base-priority</i><br><b>no priority</b>                                                                                                                                                                                               |
| <b>Context</b>     | config>service>ies>if>ipv6>vrrp                                                                                                                                                                                                                          |
| <b>Description</b> | The priority command provides the ability to configure a specific priority value to the virtual router instance. In conjunction with an optional policy command, the base-priority is used to derive the in-use priority of the virtual router instance. |

The `priority` command is only available in the non-owner vrrp virtual-router-id nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed. For non-owner virtual router instances, if the `priority` command is not executed, the base-priority will be set to 100.

The **no** form of this command restores the default value of 100 to base-priority.

**Parameters** *base-priority* — The base-priority parameter configures the base priority used by the virtual router instance. If a VRRP Priority Control policy is not also defined, the base-priority will be the in-use priority for the virtual router instance.

**Values** 1 to 254

**Default** 100

## standby-forwarding

**Syntax** **[no] standby-forwarding**

**Context** config>service>ies>if>ipv6>vrrp

**Description** This command allows the forwarding of packets by a standby router.

The **no** form of the command specifies that a standby router should not forward traffic sent to virtual router's MAC address. However, the standby router should forward traffic sent to the standby router's real MAC address.

**Default** no standby-forwarding

## telnet-reply

**Syntax** **[no] telnet-reply**

**Context** config>service>ies>if>ipv6>vrrp

**Description** This command enables the non-owner master to reply to TCP port 23 Telnet requests directed at the virtual router instances IP addresses. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Proper login and CLI command authentication is still enforced.

When telnet-reply is not enabled, TCP port 23 Telnet packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to Telnet requests regardless of the telnet-reply configuration.

The **telnet-reply** command is only available in non-owner VRRP nodal context. If the telnet-reply command is not executed, Telnet packets to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all Telnet packets destined to the non-owner virtual router instance IP addresses.

**Default** no telnet-reply

## traceroute-reply

**Syntax** [no] traceroute-reply

**Context** config>service>ies>if>ipv6>vrrp

**Description** This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.

When this command is enabled, a non-owner master can reply to traceroute requests directed to the virtual router instance IP addresses.

A non-owner backup virtual router never responds to such traceroute requests regardless of the **trace-route-reply** status.

**Default** no traceroute-reply

## 2.5.2.11 IES Spoke SDP Commands

### spoke-sdp

**Syntax** spoke-sdp sdp-id[:vc-id] [vc-type {ether | ipipe}] [create]  
no spoke-sdp sdp-id[:vc-id]

**Context** config>service>ies>if  
config>service>ies>redundant-interface

**Description** This command binds a service to an existing Service Distribution Point (SDP).

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

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with an IES service. If the **sdp sdp-id** is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router. The spoke SDP must be shut down first before it can be deleted from the configuration.

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>       | No <i>sdp-id</i> is bound to a service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Special Cases</b> | <b>IES</b> — At most, only one <i>sdp-id</i> can be bound to an IES service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>    | <p><i>sdp-id</i> — The SDP identifier. Allowed values are integers in the range of 1 and 17407 for existing SDPs.</p> <p><i>vc-id</i> — The virtual circuit identifier.</p> <p><b>Values</b> 1 to 4294967295</p> <p><i>vc-type</i> — The encapsulation and pseudowire type for the spoke-sdp.</p> <p><b>Values</b> ether: specifies Ethernet pseudowire as the type of virtual circuit (VC) associated with the SDP binding<br/>         ipipe: specifies lpipe pseudowire as the type of virtual circuit (VC) associated with the SDP binding</p> <p><b>Default</b> ether</p> |

## egress

|                    |                                                                        |
|--------------------|------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>egress</b>                                                          |
| <b>Context</b>     | config>service>ies>if>spoke-sdp<br>config>service>ies>red-if>spoke-sdp |
| <b>Description</b> | This command configures the egress SDP context.                        |

## qos

|                |                                                                                                                                                                             |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>qos</b> <i>network-policy-id</i> <b>port-redirect-group</b> <i>queue-group-name</i> [ <b>instance</b> <i>instance-id</i> ]<br><b>no qos</b> [ <i>network-policy-id</i> ] |
| <b>Context</b> | config>service>pw-template>egress<br>config>service>vprn>if>spoke-sdp>egress<br>config>service>ies>if>spoke-sdp>egress<br>config>service>ies>if>spoke-sdp>ingress           |

---

**Description** This command is used to redirect pseudowire packets to an egress port queue-group for the purpose of shaping.

The egress pseudowire shaping provisioning model allows the mapping of one or more pseudowires to the same instance of queues, or policers and queues, which are defined in the queue-group template.

Operationally, the provisioning model consists of the following steps:

1. Create an egress queue-group template and configure queues only or policers and queues for each FC that needs to be redirected.
2. Apply the queue-group template to the network egress context of all ports where there exists a network IP interface on which the pseudowire packets can be forwarded. This creates one instance of the template on the egress of the port. One or more instances of the same template can be created.
3. Configure FC-to-policer or FC-to-queue mappings together with the redirect to a queue-group in the egress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different pseudowires to different queue-group templates.
4. Apply this network QoS policy to the egress context of a spoke-SDP inside a service or to the egress context of a pseudowire template and specify the redirect queue-group name.

One or more spoke-SDPs can have their FCs redirected to use queues only or queues and policers in the same queue-group instance.

The following are the constraints and rules of this provisioning model:

1. When a pseudowire FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name does not exist, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface on which the pseudowire packet is forwarded. This queue can be a queue-group queue, or the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port. This is the existing implementation and default behavior for a pseudowire packet.
2. When a pseudowire FC is redirected to use a queue or a policer, and a queue in a queue-group and the queue-group name exists, but the policer-id and/or the queue-id is not defined in the queue-group template, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the pseudowire packet is forwarded on.
3. When a pseudowire FC is redirected to use a queue, or a policer and a queue in a queue-group, and the queue-group name exists and the policer-id or policer-id plus queue-id exist, it is not required to check that an instance of that queue-group exists in all egress network ports which have network IP interfaces. The handling of this is dealt with in the data path as follows:

When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name exists on that egress port, the packet is processed by the queue-group policer and will then be fed to the queue-group queue.

When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name does not exist on that egress port, the pseudowire packet will be fed directly to the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

- If a network QoS policy is applied to the egress context of a pseudowire, any pseudowire FC, which is not explicitly redirected in the network QoS policy, will have the corresponding packets feed directly the corresponding the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

When the queue-group name the pseudowire is redirected to exists and the redirection succeeds, the marking of the packet DEI/dot1p/DSCP and the tunnel DEI/dot1p/DSCP/EXP is performed; according to the relevant mappings of the (FC, profile) in the egress context of the network QoS policy applied to the pseudowire. This is true regardless, whether an instance of the queue-group exists or not on the egress port to which the pseudowire packet is forwarded. If the packet profile value changed due to egress child policer CIR profiling, the new profile value is used to mark the packet DEI/dot1p and the tunnel DEI/dot1p/EXP, and the DSCP/prec will be remarked if **enable-dscp-prec-marking** is enabled under the policer.

When the queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the marking of the packet DEI/dot1p/DSCP and the tunnel DEI/dot1p/DSCP/EXP fields is performed according to the relevant commands in the egress context of the network QoS policy applied to the network IP interface to which the pseudowire packet is forwarded.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

**Parameters** *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

**Values** 1 to 65535

*queue-group-name* — This optional parameter specifies that the *queue-group-name* will be used for all egress forwarding class redirections within the network QoS policy ID. The specified *queue-group-name* must exist as a port egress queue group on the port associated with the IP interface.

*instance-id* — Specifies the identification of a specific instance of the queue-group.

**Values** 1 to 16384

## vc-label

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

**Context** config>service>ies>if>spoke-sdp>egress



```
config>service>ies>red-if>spoke-sdp>egress
```

**Description** This command configures the static MPLS VC label used by this device to send packets to the far-end device in this service via this SDP.

**Parameters** *egress-vc-label* — A VC egress value that indicates a specific connection.

**Values** 16 to 1048575

## entropy-label

**Syntax** [no] **entropy-label**

**Context** config>service>ies>if>spoke-sdp

**Description** This command enables or disables the use of entropy labels on a spoke-SDP bound to an IES interface.

If **entropy-label** is configured, the entropy label and ELI are inserted in packets for which at least one LSP in the stack for the far-end of the tunnel used by the service has advertised entropy-label-capability. If the tunnel is RSVP, **entropy-label** can also be controlled under the **config>router>mpls** or **config>router>mpls>lsp** contexts.

The entropy label and hash label features are mutually exclusive. The entropy label cannot be configured on a spoke-sdp or service where the hash label feature has already been configured.

**Default** no entropy-label

## hash-label

**Syntax** **hash-label [signal-capability]**  
**no hash-label**

**Context** config>service>ies>if>spoke-sdp

**Description** This command enables the use of the hash label on a VLL, VPLS, or VPRN service bound to any MPLS-type encapsulated SDP, as well as to a VPRN service using **auto-bind-tunnel** with the **resolution-filter** configured as any MPLS tunnel type. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the gre option.

When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to 1 to indicate that.

In order to allow for applications whereby 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 will always be in the range [524,288 to 1,048,575] and will not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label will not match a value in the reserved label range.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. For VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL pseudowire packets.

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

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

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

- The local PE will insert the flow label interface parameters sub-TLV with F=1 in the PW ID FEC element in the label mapping message for that spoke-sdp or mesh-sdp.
- If the remote PE includes this sub-TLV with F=1 or F=0, then local PE must insert the hash label in the user and control plane packets.
- If remote PE does not include this sub-TLV (for example, it does not support it, or it is supported but the user did not enable the **hash-label** option or the **signal-capability** option), then the local PE establishes the PW but must not insert the hash label in the user and control packets over that spoke-sdp or mesh-sdp. If the remote PE does not support the **signal-capability** option, then there are a couple of possible outcomes:
  - If the **hash-label** option was enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the PW packets received by the local PE will have the hash label included. These packets must be dropped. The only way to solve this is to disable the signaling capability option on the local node which will result in the insertion of the hash label by both PE nodes.
  - If the **hash-label** option is not supported or was not enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the PW received by the local PE will not have the hash label included.
- The user can enable or disable the signal-capability option in CLI as needed. When doing so, the router must withdraw the label it sent to its peer and send a new label mapping message with the new value of the F bit in the flow label interface parameters sub-TLV of the PW ID FEC element.

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

---

|                   |                                                                                                                                                                                                                |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no hash-label                                                                                                                                                                                                  |
| <b>Parameters</b> | <b>signal-capability</b> — Enables the signaling and negotiation of the use of the hash label between the local and remote PE nodes. The <b>signal-capability</b> option is not supported on a VPRN spoke-sdp. |

## ingress

|                    |                                                                               |
|--------------------|-------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ingress</b>                                                                |
| <b>Context</b>     | config>service>ies>if>spoke-sdp<br>config>service>ies>red-if>spoke-sdp>egress |
| <b>Description</b> | This command configures the ingress SDP context.                              |

## qos

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>qos</b> <i>network-policy-id</i> <b>fp-redirect-group</b> <i>queue-group-name</i> <b>instance</b> <i>instance-id</i><br><b>no qos</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>service>pw-template>ingress<br>config>service>vprn>if>spoke-sdp>ingress<br>config>service>ies>if>spoke-sdp>ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command is used to redirect pseudowire packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.</p> <p>The ingress pseudowire rate-limiting feature uses a policer in queue-group provisioning model. This model allows the mapping of one or more pseudowires to the same instance of policers, which are defined in a queue-group template.</p> <p>Operationally, the provisioning model in the case of the ingress pseudowire shaping feature consists of the following steps:</p> <ol style="list-style-type: none"> <li><b>Step 1.</b> Create an ingress queue-group template and configure policers for each FC that needs to be redirected and optionally, for each traffic type (unicast, broadcast, unknown, or multicast).</li> <li><b>Step 2.</b> Apply the queue-group template to the network ingress forwarding plane where there exists a network IP interface to which the pseudowire packets can be received. This creates one instance of the template on the ingress of the FP. One or more instances of the same template can be created.</li> <li><b>Step 3.</b> Configure FC-to-policer mappings together with the policer redirect to a queue-group in the ingress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different pseudowires to different queue-group templates.</li> </ol> |

- Step 4.** Apply this network QoS policy to the ingress context of a spoke-SDP inside a service, or to the ingress context of a pseudowire template, and specify the redirect queue-group name.
- Step 5.** One or more spoke-SDPs can have their FCs redirected to use policers in the same policer queue-group instance.

The following are the constraints and rules of this provisioning model when used in the ingress pseudowire rate-limiting feature:

- Step 1.** When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name does not exist, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
- Step 2.** When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists but the policer-id is not defined in the queue-group template, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
- Step 3.** When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists and the policer-id is defined in the queue-group template, it is not required to check that an instance of that queue-group exists in all ingress FPs which have network IP interfaces. The handling of this is dealt with in the data path as follows:

When a pseudowire packet for that FC is received and an instance of the referenced queue-group name exists on that FP, the packet is processed by the policer and will then feed the per-FP ingress shared queues referred to as *policer-output-queues*.

When a pseudowire packet for that FC is received and an instance of the referenced queue-group name does not exist on that FP, the pseudowire packets will be fed directly into the corresponding ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.

- If a network QoS policy is applied to the ingress context of a pseudowire, any pseudowire FC which is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
- If no network QoS policy is applied to the ingress context of the pseudowire, then all packets of the pseudowire will feed:

the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the MDA/FP. This is the default behavior.

a queue-group policer followed by the per-FP ingress shared queues referred to as *policer-output-queues* if the ingress context of the network IP interface from which the packet is received is redirected to a queue-group (csc-policing). The only exceptions to this behavior are for packets received from a IES/VPRN spoke interface and from an R-VPLS spoke-SDP, which is forwarded to the R-VPLS IP interface. In these two cases, the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the MDA/FP is used.

When a pseudowire is redirected to use a policer queue-group, the classification of the packet for the purpose of FC and profile determination is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the pseudowire. This is true regardless of whether an instance of the named policer queue-group exists on the ingress FP on which the pseudowire packet is received. The user can apply a QoS filter matching the dot1.p in the VLAN tag corresponding to the Ethernet port encapsulation, the EXP in the outer label when the tunnel is an LSP, the DSCP in the IP header if the tunnel encapsulation is GRE, and the DSCP in the payload IP header if the user enabled the **ler-use-dscp** option and the pseudowire terminates in IES or VPRN service (spoke-interface).

When the policer queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the packet classification is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface on which the pseudowire packet is received.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

**Parameters** *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

**Values** 1 to 65535

**fp-redirect-group** *queue-group-name* — Specifies the name of the queue group template up to 32 characters in length.

**ingress-instance** *instance-id* — Specifies the identification of a specific instance of the queue-group.

**Values** 1 to 16384

## vc-label

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

**Context** config>service>ies>if>spoke-sdp>ingress  
config>service>ies>red-if>spoke-sdp>ingress

**Description** This command configures the static MPLS VC label used by the far-end device to send packets to this device in this service via this SDP.

---

|                   |                                                                                    |
|-------------------|------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>ingress-vc-label</i> — A VC ingress value that indicates a specific connection. |
| <b>Values</b>     | 2048 to 18431                                                                      |

## accounting-policy

|                    |                                                                               |
|--------------------|-------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>accounting-policy</b> <i>acct-policy-id</i><br><b>no accounting-policy</b> |
| <b>Context</b>     | config>service>ies>if>spoke-sdp                                               |
| <b>Description</b> | This command configures an accounting-policy.                                 |
| <b>Parameters</b>  | <i>acct-policy-id</i> — Specifies an accounting policy ID.                    |
| <b>Values</b>      | 1 to 99                                                                       |

## app-profile

|                    |                                                                     |
|--------------------|---------------------------------------------------------------------|
| <b>Syntax</b>      | <b>app-profile</b> <i>app-profile-name</i><br><b>no app-profile</b> |
| <b>Context</b>     | config>service>ies>if>spoke-sdp                                     |
| <b>Description</b> | This command configures the application profile name.               |
| <b>Parameters</b>  | <i>app-profile-name</i> — Specifies the application profile name.   |

## bfd-enable

|                    |                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>bdf-enable</b><br><b>no bfd-enable</b>                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies>if>spoke-sdp                                                                                                                                                                                                                      |
| <b>Description</b> | This command enables VCCV BFD on the PW associated with the VLL, BGP VPWS, or VPLS service. The parameters for the BFD session are derived from the named BFD template, which must have been first configured using the <b>bfd-template</b> command. |

## bfd-template

|                |                                                           |
|----------------|-----------------------------------------------------------|
| <b>Syntax</b>  | <b>bdf-template</b> <i>name</i><br><b>no bfd-template</b> |
| <b>Context</b> | config>service>ies>if>spoke-sdp                           |

---

|                    |                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command configures a named BFD template to be used by VCCV BFD on PWs belonging to the VLL, BGP VPWS, or VPLS service. The template specifies parameters, such as the minimum transmit and receive control packet timer intervals, to be used by the BFD session. Template parameters are configured under the <b>config&gt;router&gt;bfd</b> context. |
| <b>Default</b>     | no bfd-template                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <i>name</i> — A text string name for the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.                                                                                                                                                                                                                               |

## collect-stats

|                    |                                                         |
|--------------------|---------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] collect-stats</b>                               |
| <b>Context</b>     | config>service>ies>if>spoke-sdp                         |
| <b>Description</b> | This command enables or disables statistics collection. |

## control-channel-status

|                    |                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>control-channel-status</b>                                                                                                              |
| <b>Context</b>     | config>service>ies>if>spoke-sdp                                                                                                            |
| <b>Description</b> | This command enables the configuration of static pseudowire status signaling on a spoke-SDP for which signaling for its SDP is set to OFF. |

A control-channel-status no shutdown is allowed only if all of the following are true:

- SDP signaling is off.
- The control-word is enabled (the control-word is disabled by default)
- The service type is Epipe, Apipe, VPLS, Cpipe, or IES/VPRN
- Mate SDP signaling is off (in vc-switched services)
- The pw-path-id is configured for this spoke-SDP.

The **no** form of this command removes control channel status signaling from a spoke-SDP. It can only be removed if control channel status is shut down.

|                |                           |
|----------------|---------------------------|
| <b>Default</b> | no control-channel-status |
|----------------|---------------------------|

## acknowledgment

|                |                                                        |
|----------------|--------------------------------------------------------|
| <b>Syntax</b>  | <b>[no] acknowledgment</b>                             |
| <b>Context</b> | config>service>ies>if>spoke-sdp>control-channel-status |

---

**Description** This command enables the acknowledgment of control channel status messages. By default, no acknowledgment packets are sent.

## refresh-timer

**Syntax** **refresh-timer** *value*  
**no refresh-timer**

**Context** config>service>ies>if>spoke-sdp>control-channel-status

**Description** This command configures the refresh timer for control channel status signaling packets. By default, no refresh packets are sent.

**Default** no refresh-timer

**Parameters** *value* — Specifies the refresh timer value.

**Values** 10 to 65535 seconds

**Default** 0 (off)

## request-timer

**Syntax** **request-timer** *timer1* **retry-timer** *timer2* **timeout-multiplier** *multiplier*  
**no request-timer**

**Context** config>service>ies>if>spoke-sdp>control-channel-status

**Description** This command configures the control channel status request mechanism. When it is configured, control channel status request procedures are used. These augment the procedures for control channel status messaging from RFC 6478. This command is mutually exclusive with a non-zero refresh-timer value.

**Parameters** *timer1* — Specifies the interval at which pseudowire status messages, including a reliable delivery TLV, with the “request” bit set, are sent.

**Values** 10 to 65535 seconds

**retry-timer** *timer2* — Specifies the timeout interval if no response to a pseudowire status request is received. This parameter must be configured. A value of zero (0) disables retries.

**Values** 0, 3 to 60 seconds

**timeout-multiplier** *multiplier* — If a requesting node does not receive a valid response to a pseudowire status request within this multiplier times the retry timer, then it will assume the pseudowire is down. This parameter is optional.

**Values** 3 to 20 seconds



## control-word

|                    |                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] control-word</b>                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>service>ies>if>spoke-sdp                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command enables/disables the PW control word on spoke-sdps terminated on an IES or VPRN interface. The control word must be enabled to allow MPLS-TP OAM on the spoke-sdp</p> <p>It is only valid for MPLS-TP spoke-sdps when used with IES and VPRN services.</p> |
| <b>Default</b>     | no control-word                                                                                                                                                                                                                                                            |

### 2.5.2.12 IES SAP Commands

## sap

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sap sap-id [create]</b><br><b>no sap sap-id</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>service>ies>if<br>config>service>ies>sub-if>grp-if                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique.</p> <p>All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.</p> <p>Enter an existing SAP without the <b>create</b> keyword to edit SAP parameters. The SAP is owned by the service in which it was created.</p> <p>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 <b>config interface port-type port-id mode access</b> command. For the 7750 SR, channelized TDM ports are always access ports.</p> <p>If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.</p> <p>You can configure an IES interface as a loopback interface by issuing the <b>loopback</b> command instead of the <b>sap sap-id</b> 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.</p> |

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP will also be deleted. For Internet Enhanced Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed. The no form of this command causes the ptp-hw-assist to be disabled.

**Default** No SAPs are defined.

**Special Cases** **IES** — For the 7750 SR, an IES SAP can be defined with Ethernet ports, SONET/SDH or TDM channels. For the 7450 ESS, IES SAP can be defined with Ethernet or SONET/SDH ports. A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. For the 7750 SR, group interfaces allow more than one SAP. Attempts to create a second SAP on an IP interface fails and generate an error; the original SAP will not be affected.

Command syntax for the 7750 SR: **sap ipsec-id.private | public:tag** associates an IPSec group SAP with this interface. This is the public side for an IPSec tunnel. Tunnels referencing this IPSec group in the private side may be created if their local IP is in the subnet of the interface subnet and the routing context specified matches with the one of the interface.

This context will provide a SAP to the tunnel. The operator may associate an ingress and egress QoS policies as well as filters and virtual scheduling contexts. Internally this creates an Ethernet SAP that will be used to send and receive encrypted traffic to and from the MDA. Multiple tunnels can be associated with this SAP. The “tag” will be a dot1q value. The operator may see it as an identifier. The range is limited to 1 to 4095.

**Parameters** *sap-id* — Specifies the physical port identifier portion of the SAP definition.

*port-id* — Specifies the physical port ID.

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. When the *port-id* parameter represents SONET/SDH and TDM channels (7750 SR), the port ID must include the channel ID. A period “.” separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

|                   |                                 |         |         |
|-------------------|---------------------------------|---------|---------|
| <i>port-id</i>    | <i>slot/mda/port [.channel]</i> |         |         |
| <i>eth-sat-id</i> | <i>esat-id/slot/port</i>        |         |         |
|                   | <i>esat</i>                     | keyword |         |
|                   | <i>id</i>                       |         | 1 to 20 |
| <i>pxc-id</i>     | <i>pxc-id.sub-port</i>          |         |         |
|                   | <i>pxc</i>                      | keyword |         |
|                   | <i>id</i>                       |         | 1 to 64 |
|                   | <i>sub-port</i>                 |         | a, b    |

**create** — Keyword used to create a SAP instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

## aarp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>aarp</b> <i>aarpId</i> <b>type</b> <i>type</i><br><b>no aarp</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies>if>sap<br>config>service>ies>if>spoke-sdp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command associates an AARP instance with a multi-homed SAP or spoke SDP. This instance uses the same AARP ID in the same node or in a peer node (pre-configured) to provide traffic flow and packet asymmetry removal for a multi-homed SAP or spoke SDP.</p> <p>The type specifies the role of this service point in the AARP: either, primary (dual-homed) or secondary (dual-homed-secondary). The AA service attributes (app-profile and transit-policy) of the primary are inherited by the secondary endpoints. All endpoints within an AARP must be of the same type (SAP or spoke), and all endpoints with an AARP must be within the same service.</p> <p>The no form of the command removes the association between an AARP instance and a multi-homed SAP or spoke SDP.</p> |
| <b>Default</b>     | no aarp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><i>aarpId</i> — Specifies the AARP instance associated with this SAP. If not configured, no AARP instance is associated with this SAP.</p> <p><b>Values</b> 1 to 65535</p> <p><i>type</i> — Specifies the role of the SAP referenced by the AARP instance.</p> <p><b>Values</b> <b>dual-homed</b> — The primary dual-homed AA subscriber side service-point of an AARP instance; only supported for Epipe, IES, and VPRN SAP and spoke SDP.</p> <p><b>dual-homed-secondary</b> — One of the secondary dual-homed AA subscriber side service-points of an AARP instance; only supported for Epipe, IES, and VPRN SAP and spoke SDP.</p>                                                                                                                                                      |

## anti-spoof

|                |                                                               |
|----------------|---------------------------------------------------------------|
| <b>Syntax</b>  | <b>anti-spoof</b> {ip   mac   ip-mac}<br><b>no anti-spoof</b> |
| <b>Context</b> | config>service>ies>if>sap                                     |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP.</p> <p>The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (<b>ip</b>, <b>ip-mac</b>, <b>nh-mac</b>) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.</p> <p>The <b>no</b> form of the command disables anti-spoof filtering on the SAP.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>     | no anti-spoof                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <p><b>ip</b> — Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the anti-spoof type <b>ip</b> command fails.</p> <p><b>mac</b> — Configures SAP anti-spoof filtering to use only the source MAC address in its lookup. Setting the anti-spoof filter type to <b>mac</b> is not allowed on non-Ethernet encapsulated SAPs. If a static host exists on the SAP without a specified MAC address, the anti-spoof type <b>mac</b> command fails. The anti-spoof type <b>mac</b> command will also fail if the SAP does not support Ethernet encapsulation.</p> <p><b>ip-mac</b> — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup. If a static host exists on the SAP without both the IP address and MAC address specified, the anti-spoof type <b>ip-mac</b> command fails. This is also true if the default anti-spoof filter type of the SAP is <b>ip-mac</b> and the default is not overridden. The anti-spoof type <b>ip-mac</b> command will also fail if the SAP does not support Ethernet encapsulation.</p> |

## anti-spoof

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>anti-spoof {ip   ip-mac   nh-mac}</b><br><b>no anti-spoof</b>                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP.</p> <p>The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (<b>ip</b>, <b>ip-mac</b>) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.</p> <p>The <b>no</b> form of the command reverts to the default.</p> |
| <b>Default</b>     | anti-spoof ip-mac                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

- Parameters**
- ip** — Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the anti-spoof type **ip** command fails.
  - ip-mac** — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup. If a static host exists on the SAP without both the IP address and MAC address specified, the anti-spoof type **ip-mac** command fails. This is also true if the default anti-spoof filter type of the SAP is **ip-mac** and the default is not overridden. The anti-spoof type **ip-mac** command will also fail if the SAP does not support Ethernet encapsulation.
  - nh-mac** — Indicates that the ingress anti-spoof is based on the source MAC address and the egress anti-spoof is based on the nh-ip-address.

## app-profile

- Syntax** **app-profile** *app-profile-name*  
**no app-profile**
- Context** config>service>ies>if>sap  
config>service>ies>sub-if>grp-if>sap
- Description** This command configures the application profile name.
- Parameters** *app-profile-name* — Specifies an existing application profile name configured in the **config>app-assure>group>policy** context.

## ip-tunnel

- Syntax** **ip-tunnel** *name* [**create**]  
**no ip-tunnel** *name*
- Context** config>service>ies>if>sap
- Description** This command is used to configure an IP-GRE or IP-IP tunnel and associate it with a private tunnel SAP within an IES or VPRN service.
- The **no** form of the command deletes the specified IP/GRE or IP-IP tunnel from the configuration. The tunnel must be administratively shutdown before issuing the no ip-tunnel command.
- Default** no-ip tunnel *name*
- Parameters** **ip-tunnel** *name* — Specifies the name of the IP tunnel. Tunnel 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.

## lag-link-map-profile

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lag-link-map-profile</b> <i>lag-link-map-profile-id</i><br><b>no lag-link-map-profile</b>                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>service>ies>if>sap<br>config>service>ies>sub-if>grp-if>sap                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command assigns a pre-configured lag link map profile to a SAP/network interface configured on a LAG or a PW port that exists on a LAG. Once assigned/de-assigned, the SAP/network interface egress traffic will be re-hashed over LAG as required by the new configuration.<br><br>The <b>no</b> form of this command reverts the SAP/network interface to use per-flow, service or link hash as configured for the service/LAG. |
| <b>Default</b>     | no lag-link-map-profile                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <i>lag-link-map-profile-id</i> — An integer from 1 to 64 that defines a unique lag link map profile on which the LAG the SAP/network interface exist.                                                                                                                                                                                                                                                                                  |

## lag-per-link-hash

|                    |                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lag-per-link-hash class</b> {1   2   3} <b>weight</b> <i>weight</i><br><b>no per-link-hash</b>                                                                                                 |
| <b>Context</b>     | config>service>ies>if>sap<br>config>service>ies>sub-if>grp-if>sap                                                                                                                                 |
| <b>Description</b> | This command configures weight and class to this SAP to be used on LAG egress when the LAG uses weighted per-link-hash.<br><br>The <b>no</b> form of this command restores default configuration. |
| <b>Default</b>     | no lag-per-link-hash (equivalent to weight 1 class 1)                                                                                                                                             |
| <b>Parameters</b>  | <i>weight</i> — Specifies the weight.<br><br><b>Values</b> 1 to 1024                                                                                                                              |

## multi-service-site

|                |                                                                                                               |
|----------------|---------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>multi-service-site</b> <i>customer-site-name</i><br><b>no multi-service-site</b> <i>customer-site-name</i> |
| <b>Context</b> | config>service>ies>if>sap<br>config>service>ies>sub-if>grp-if>sap                                             |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command creates a new customer site or edits an existing customer site with the <i>customer-site-name</i> parameter. A customer site is an anchor point to create an ingress and egress virtual scheduler hierarchy. On the 7750 SR, when a site is created, it must be assigned to a chassis slot or port. When a site is created, it must be assigned to a chassis slot or port with the exception of the 7450 ESS-1 in which the slot is set to 1. When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).</p> <p>The scheduler policy association with the customer site normally prevents the scheduler policy from being deleted until after the scheduler policy is removed from the customer site. The multi-service-site object will generate a log message indicating that the association was deleted due to scheduler policy removal.</p> <p>When the multi-service customer site is created, an ingress and egress scheduler policy association does not exist. This does not prevent the site from being assigned to a chassis slot or prevent service SAP assignment. After the site has been created, the ingress and egress scheduler policy associations can be assigned or removed at any time.</p>                                                                                                                                                                                                                                                     |
| <b>Default</b>     | n/a — Each customer site must be explicitly created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <p><i>customer-site-name</i> — Each customer site must have a unique name within the context of the customer. If <i>customer-site-name</i> already exists for the customer ID, the CLI context changes to that site name for the purpose of editing the site scheduler policies or assignment. Any modifications made to an existing site will affect all SAPs associated with the site. Changing a scheduler policy association may cause new schedulers to be created and existing policers and queues on the SAPs to no longer be orphaned. Existing schedulers on the site may cease to exist, causing policers and queues relying on that scheduler to be orphaned.</p> <p>If the <i>customer-site-name</i> does not exist, it is assumed that an attempt is being made to create a site of that name in the customer ID context. The success of the command execution depends on the following:</p> <ul style="list-style-type: none"> <li>The maximum number of customer sites defined for the chassis has not been met.</li> <li>The <i>customer-site-name</i> is valid.</li> <li>The <b>create</b> keyword is included in the command line syntax (if the system requires it).</li> </ul> <p>When the maximum number of customer sites has been exceeded a configuration error occurs; the command will not execute and the CLI context will not change.</p> <p>If the <i>customer-site-name</i> is invalid, a syntax error occurs; the command will not execute and the CLI context will not change.</p> <p><b>Values</b> Valid names consist of any string up to 32 characters long 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.</p> |

---

## static-host

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>static-host ip</b> <i>ip/did-address</i> [ <b>mac</b> <i>ieee-address</i> ] [ <b>create</b> ]<br><b>static-host mac</b> <i>ieee-address</i> [ <b>create</b> ]<br><b>no static-host</b> [ <b>ip</b> <i>ip-address</i> >] <b>mac</b> <i>ieee-address</i> ><br><b>no static-host all</b> [ <b>force</b> ]<br><b>no static-host ip</b> <i>ip-address</i>                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>service>ies>if>sap<br>config>service>ies>sub-if>grp-if>sap                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | This command configures a static host on this SAP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <b>ip</b> <i>ip-address</i> — Specifies the IPv4 unicast address.<br><br><b>mac</b> <i>ieee-address</i> — Specify this optional parameter when defining a static host. Every static host definition must have at least one address defined, IP or MAC.<br><br><b>force</b> — Specifies the forced removal of the static host addresses.<br><br><b>sla-profile</b> <i>sla-profile-name</i> — This optional parameter is used to specify an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the <b>config&gt;subscr-mgmt&gt;sla-profile</b> context. |

## ancp-string

|                    |                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ancp-string</b> <i>ancp-string</i><br><b>no ancp-string</b>                            |
| <b>Context</b>     | config>service>ies>if>sap>static-host<br>config>service>ies>sub-if>grp-if>sap>static-host |
| <b>Description</b> | This command specifies the ANCP string associated to this SAP host.                       |
| <b>Parameters</b>  | <i>ancp-string</i> — Specifies the ANCP string up to 63 characters in length.             |

## app-profile

|                    |                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>app-profile</b> <i>app-profile-name</i><br><b>no app-profile</b>                             |
| <b>Context</b>     | config>service>ies>if>sap>static-host<br>config>service>ies>sub-if>grp-if>sap>static-host       |
| <b>Description</b> | This command specifies an application profile name.                                             |
| <b>Parameters</b>  | <i>app-profile-name</i> — Specifies the application profile name up to 32 characters in length. |



## inter-dest-id

|                    |                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>inter-dest-id</b> <i>intermediate-destination-id</i><br><b>no inter-dest-id</b>                                     |
| <b>Context</b>     | config>service>ies>if>sap>static-host<br>config>service>ies>sub-if>grp-if>sap>static-host                              |
| <b>Description</b> | Specifies to which intermediate destination (for example, a DSLAM) this host belongs.                                  |
| <b>Parameters</b>  | <i>intermediate-destination-id</i> — Specifies the intermediate destination identifier, up to 32 characters in length. |

## managed-routes

|                    |                                                                 |
|--------------------|-----------------------------------------------------------------|
| <b>Syntax</b>      | <b>managed-routes</b>                                           |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap>static-host>managed-routes |
| <b>Description</b> | This command configures managed routes.                         |

## route

|                    |                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>route</b> { <i>ip-prefix/length</i>   <i>ip-prefix netmask</i> } [ <b>create</b> ]<br><b>no route</b> { <i>ip-prefix/length</i>   <i>ip-prefix netmask</i> }                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap>static-host>managed-routes                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | This command assigns managed-route to a given subscriber-host. As a consequence, a static route pointing subscriber-host ip address as a next hop will be installed in FIB. Up to 16 managed routes per subscriber-host can be configured.<br><br>The <b>no</b> form of the command removes the respective route. Per default, there are no managed-routes configured. |

## sla-profile

|                    |                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sla-profile</b> <i>sla-profile-name</i><br><b>no sla-profile</b>                                                                                                                             |
| <b>Context</b>     | config>service>ies>if>sap>static-host<br>config>service>ies>sub-if>grp-if>sap>static-host                                                                                                       |
| <b>Description</b> | This command specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the <b>config&gt;subscr-mgmt&gt;sla-profile</b> context. |

---

**Parameters**     *sla-profile-name* — Specifies the SLA profile name.

## sub-profile

**Syntax**     **sub-profile** *sub-profile-name*  
              **no sub-profile**

**Context**     config>service>ies>if>sap>static-host  
              config>service>ies>sub-if>grp-if>sap>static-host

**Description**     This command specifies an existing subscriber profile name to be associated with the static subscriber host.

**Parameters**     *sub-profile-name* — Specifies the sub-profile name.

## subscriber

**Syntax**     **subscriber** *sub-ident*  
              **no subscriber**

**Context**     config>service>ies>if>sap>static-host  
              config>service>ies>sub-if>grp-if>sap>static-host

**Description**     This command specifies an existing subscriber identification profile to be associated with the static subscriber host.

**Parameters**     *sub-ident* — Specifies the subscriber identification.

## subscriber-sap-id

**Syntax**     **[no] subscriber-sap-id**

**Context**     config>service>ies>if>sap>static-host  
              config>service>ies>sub-if>grp-if>sap>static-host

**Description**     This command enables using the SAP ID as subscriber id.

**Parameters**     **subscriber-sap-id** — Specifies to use the sap-id as the subscriber-id.

## pw-path-id

**Syntax**     **[no] pw-path-id**

**Context**     config>service>ies>if>spoke-sdp  
              config>service>vprn>if>spoke-sdp

**Description** This command enters the context to configure an MPLS-TP Pseudowire Path Identifier for a spoke-sdp. All elements of the PW path ID must be configured in order to enable a spoke-sdp with a PW path ID.

For an IES or VPRN spoke-sdp, the pw-path-id is only valid for Ethernet spoke-sdps.

The **pw-path-id** is only configurable if all of the following is true:

- SDP signaling is off
- control-word is enabled (control-word is disabled by default)
- the service type is epipe, vpls, cpipe, apipe, or IES/VPRN interface
- mate SDP signaling is off for vc-switched services

The **no** form of the command deletes the PW path ID.

**Default** no pw-path-id

## agi

**Syntax** **agi** *agi*  
**no agi**

**Context** config>service>ies>if>spoke-sdp>pw-path-id  
config>service>vprn>if>spoke-sdp>pw-path-id

**Description** This command configures the attachment group identifier for an MPLS-TP PW.

**Parameters** *agi* — Specifies the attachment group identifier.

**Values** 0 to 4294967295

## saii-type2

**Syntax** **saii-type2** *global-id:node-id:ac-id*  
**no saii-type2**

**Context** config>service>ies>if>spoke-sdp>pw-path-id  
config>service>vprn>if>spoke-sdp>pw-path-id

**Description** This command configures the source individual attachment identifier (SAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured (for example, it is at an S-PE), then the values must match those of the taii-type2 of the mate spoke-sdp.

**Parameters** *global-id* — Specifies the global ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.

**Values** 0 to 4294967295

*node-id* — Specifies the node ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.

**Values** a.b.c.d or 0 to 4294967295

*ac-id* — Specifies the attachment circuit ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value.

**Values** 1 to 4294967295

## taii-type2

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>taii-type2</b> <i>global-id:node-id:ac-id</i><br><b>no taii-type2</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>service>ies>if>spoke-sdp>pw-path-id<br>config>service>vprn>if>spoke-sdp>pw-path-id                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | This command configures the target individual attachment identifier (TAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured (for example, it is at an S-PE), then the values must match those of the saii-type2 of the mate spoke-sdp.                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <p><i>global-id</i> — Specifies the global ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.</p> <p><b>Values</b> 0 to 4294967295</p> <p><i>node-id</i> — Specifies the node ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.</p> <p><b>Values</b> a.b.c.d or 0 to 4294967295</p> <p><i>ac-id</i> — Specifies the attachment circuit ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value.</p> <p><b>Values</b> 1 to 4294967295</p> |

## transit-policy

|                    |                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>transit-policy</b> { <b>ip</b> <i>ip-aasub-policy-id</i>   <b>prefix</b> <i>prefix-aasub-policy-id</i> }<br><b>no transit-policy</b>                                                                                                      |
| <b>Context</b>     | config>service>ies>if>sap><br>config>service>ies>if>spoke-sdp>                                                                                                                                                                               |
| <b>Description</b> | This command associates an AA transit policy to the service. The transit IP policy must be defined prior to associating the policy with a SAP in the <b>config&gt;application assurance&gt;group&gt;policy&gt;transit-ip-policy</b> context. |

Transit AA subscribers are managed by the system through this service policy, which determines how transit subs are created and removed for that service.

The no form of the command removes the association of the policy to the service.

|                   |                                                                                                  |
|-------------------|--------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no transit-policy                                                                                |
| <b>Parameters</b> | <i>ip-aasub-policy-id</i> — Specifies an integer identifying an IP transit IP profile entry.     |
| <b>Values</b>     | 1 to 65535                                                                                       |
|                   | <i>prefix-aasub-policy-id</i> — Specifies an integer identifying a prefix transit profile entry. |
| <b>Values</b>     | 1 to 65535                                                                                       |

## dynamic-tunnel-redundant-next-hop

|                    |                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dynamic-tunnel-redundant-next-hop</b> <i>ip-address</i><br><b>no dynamic-tunnel-redundant-next-hop</b>                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>if                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | This command specifies redundant next-hop address on public or private IPSec interface (with public or private tunnel-sap) for dynamic IPSec tunnel. The specified next-hop address will be used by standby node to shunt traffic to master in case of it receives them.<br><br>The next-hop address will be resolved in routing table of corresponding service. |
| <b>Default</b>     | none                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies the dynamic ISA tunnel redundant next-hop address.                                                                                                                                                                                                                                                                                 |

## enable-ingress-stats

|                    |                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] enable-ingress-stats</b>                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies >interface<br>config>service>ies>sub-if>grp-if                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | This command enables the collection of ingress interface IP stats. This command is only applicable to IP statistics, and not to uRPF statistics.<br><br>If enabled, then the following statistics are collected: <ul style="list-style-type: none"> <li>• IPv4 offered packets</li> <li>• IPv4 offered octets</li> <li>• IPv6 offered packets</li> <li>• IPv6 offered octets</li> </ul> |

- Octet statistics for IPv4 and IPv6 bytes at IP interfaces include the layer 2 frame overhead.

**Default** no enable-ingress-stats

## egr-ip-load-balancing

**Syntax** **egr-ip-load-balancing** {**source** | **destination** | **inner-ip**}  
**no egr-ip-load-balancing**

**Context** config>service>ies>if>load-balancing

**Description** This command specifies whether to include source address or destination address or both in LAG/ECMP hash on IP interfaces. Additionally, when I4-load-balancing is enabled the command applies also to inclusion of source/destination port in the hash inputs.

The **no** form of this command includes both source and destination parameters.

**Default** no egr-ip-load-balancing

**Parameters** *source* — Specifies using source address and (if I4-load balancing is enabled) source port in the hash, ignore destination address/port.  
*destination* — Specifies using destination address and (if I4-load balancing is enabled) destination port in the hash, ignore source address/port.  
*inner-ip* — Specifies use of the inner IP header parameters instead of outer IP header parameters in LAG/ECMP hash for IPv4 encapsulated traffic.

## enable-mac-accounting

**Syntax** [**no**] **enable-mac-accounting**

**Context** config>service>ies>if

**Description** This command enables MAC accounting functionality on this interface.

The **no** form of the command disables MAC accounting functionality on this interface.

## flowspec

**Syntax** [**no**] **flowspec**

**Context** config>service>vprn>if>sap>ingress  
config>service>vprn>if>spoke-sdp>ingress  
config>service>ies>if>sap>ingress  
config>service>ies>if>spoke-sdp>ingress

**Description** This command enables IPv4 flowspec filtering on an access IP interface associated with a VPRN or IES service. Filtering is based on all of the IPv4 flowspec routes that have been received and accepted by the corresponding BGP instance. Ingress IPv4 traffic on an interface can be filtered by both a user-defined IPv4 filter and flowspec. Evaluation proceeds in this order:

- user-defined IPv4 filter entries
- flowspec-derived filter entries
- user-defined IPv4 filter default-action

The **no** form of the command removes IPv4 flowspec filtering from an IP interface.

**Default** no flowspec. No access interfaces have IPv4 flowspec enabled.

## host-connectivity-verify

**Syntax** **host-connectivity-verify** [**source** {**vrrp** | **interface**}] [**interval** *interval*] [**action** {**remove** | **alarm**}] [**timeout** *retry-timeout*] [**retry-count** *count*]  
**host-connectivity-verify** [**interval** *interval*] [**action** {**remove** | **alarm**}] [**timeout** *retry-timeout*] [**retry-count** *count*] [**family** *family*]

**Context** config>service>ies>if  
config>service>ies>sub-if>grp-if

**Description** This command enables subscriber host connectivity verification for all hosts on this interface. This tool will periodically scan all known hosts (from dhcp-state) and perform a UC ARP request. The subscriber host connectivity verification will maintain state (connected vs. not-connected) for all hosts.

**Default** no host-connectivity-verify

**Parameters** **source {interface}** — Specifies the source to be used for generation of subscriber host connectivity verification packets. The **interface** keyword forces the use of the interface mac and ip addresses. There are up to 16 possible subnets on a given interface, therefore subscriber host connectivity verification tool will use always an address of the subnet to which the given host is pertaining. In case of group-interfaces, one of the parent subscriber-interface subnets (depending on host's address) will be used.

**interval** *interval* — The interval, in minutes, which specifies the time interval which all known sources should be verified. The actual rate is then dependent on number of known hosts and interval.

**Values** 1 to 6000

A zero value can be used by the SNMP agent to disable host-connectivity-verify.

**action {remove | alarm}** — Defines the action taken on a subscriber host connectivity verification failure for a given host. The **remove** keyword raises an alarm and removes DHCP state and releases all allocated resources (queues, table entries and so on). DHCP release will be signaled to corresponding DHCP server. Static host will never be removed. The **alarm** keyword raises an alarm indicating that the host is disconnected.

**timeout *retry-timeout*** — Specifies the timeout in seconds between consecutive retries of subscriber host connectivity verification checks, in case the host does not respond.

**Values** 10 to 60 seconds

**retry-count *count*** — Specifies the number of retries that will be carried out before a subscriber host is considered to have failed the SHCV check.

**Values** 2 to 29

**family *family*** — Indicates the IP address family for which subscriber host connectivity verification checks will be enabled. It can be set to **ipv4** or **ipv6**, or **both**.

## source

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>source</b> <i>ip-address</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>service>ies>if>sap>ip-tunnel                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | This command configures the source IPv4 or IPv6 address to use for an IP tunnel. This configuration applies to the outer IP header of the encapsulated packets. The IPv4 or IPv6 address must belong to the one of the IP subnets associated with the public SAP interface of the tunnel-group. The <b>source</b> address, <b>remote-ip</b> address and <b>backup-remote-ip</b> address of a tunnel must all belong to the same address family (IPv4 or IPv6). When the source address contains an IPv6 address it must be a global unicast address. |
| <b>Default</b>     | no source                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <i>ip-address</i> — An IPv4 address or an IPv6 address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## remote-ip

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>remote-ip</b> <i>ip-address</i><br><b>no remote-ip</b>                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies>if>sap>ip-tunnel                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | This command configures the primary destination IPv4 or IPv6 address to use for an IP tunnel. This configuration applies to the outer IP header of the encapsulated packets. The <b>source</b> address, <b>remote-ip</b> address and <b>backup-remote-ip</b> address of a tunnel must all belong to the same address family (IPv4 or IPv6). When the remote-ip address contains an IPv6 address it must be a global unicast address. |



---

|                   |                                                         |
|-------------------|---------------------------------------------------------|
| <b>Default</b>    | no remote-ip                                            |
| <b>Parameters</b> | <i>ip-address</i> — An IPv4 address or an IPv6 address. |

## backup-remote-ip

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>backup-remote-ip</b> <i>ip-address</i><br><b>no backup-remote-ip</b>                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>service>ies>if>sap>ip-tunnel<br>config>service>vprn>if>sap>ip-tunnel                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | This command configures the alternate destination IPv4 or IPv6 address to use for an IP tunnel. This destination address is used only if the primary destination configured with the <b>remote-ip</b> command is unreachable in the delivery service. The <b>source</b> address, <b>remote-ip</b> address and <b>backup-remote-ip</b> address of a tunnel must all belong to the same address family (IPv4 or IPv6). When the backup-remote-ip address contains an IPv6 address it must be a global unicast address. |
| <b>Default</b>     | no backup-remote-ip                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <i>ip-address</i> — An IPv4 address or an IPv6 address.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## clear-df-bit

|                    |                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] clear-df-bit</b>                                                                                     |
| <b>Context</b>     | config>service>if>ies>sap<br>config>service>ies>if>sap>ip-tunnel                                             |
| <b>Description</b> | This command specifies whether to clear the Do not Fragment (DF) bit in the outgoing packets in this tunnel. |

## 2.5.2.13 SAP Subscriber Management Commands

### sub-sla-mgmt

|                    |                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] sub-sla-mgmt</b>                                                                    |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap                                                        |
| <b>Description</b> | This command enters the context to configure subscriber management parameters for this SAP. |

---

**Default**    no sub-sla-mgmt

## def-sla-profile

**Syntax**    **def-sla-profile** *default-sla-profile-name*  
**no def-sla-profile**

**Context**    config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt

**Description**    This command specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sla-profile** context.

An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts. SLA profiles are maintained in two locations, the subscriber identification policy and the subscriber profile templates. After a subscriber host is associated with an SLA profile name, either the subscriber identification policy used to identify the subscriber or the subscriber profile associated with the subscriber host must contain an SLA profile with that name. If both the subscriber identification policy and the subscriber profile contain the SLA profile name, the SLA profile in the subscriber profile is used.

The **no** form of the command removes the default SLA profile from the SAP configuration.

**Default**    no def-sla-profile

**Parameters**    *default-sla-profile-name* — Specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sla-profile** context.

## def-sub-profile

**Syntax**    **def-sub-profile** *default-subscriber-profile-name*

**Context**    config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt

**Description**    This command specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sub-profile** context.

A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscriber using the subscriber profile. Subscriber profiles also allow for specific SLA profile definitions when the default definitions from the subscriber identification policy must be overridden.

The **no** form of the command removes the default SLA profile from the SAP configuration.

**Parameters** *default-sub-profile* — Specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sub-profile** context.

## sub-ident-policy

**Syntax** **sub-ident-policy** *sub-ident-policy-name*

**Context** config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt

**Description** This command associates a subscriber identification policy to this SAP. The subscriber identification policy must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sub-ident-policy** context.

Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.

For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet string. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.

When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.

The **no** form of the command removes the default subscriber identification policy from the SAP configuration.

**Default** no sub-ident-policy

**Parameters** *sub-ident-policy-name* — Specifies a subscriber identification policy for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sub-ident-policy** context.

## multi-sub-sap

**Syntax** **multi-sub-sap** [*subscriber-limit*]  
**no multi-sub-sap**

**Context** config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt

**Description** This command configures the maximum number of subscribers for this SAP.

The **no** form of this command returns the default value.

**Default** multi-sub-sap 1

---

**Parameters**     *subscriber-limit* — Specifies the maximum number of subscribers for this SAP.

**Values**        2 to 8000

## single-sub-parameters

**Syntax**        **single-sub-parameters**

**Context**        config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt

**Description**    This command enters the context to configure single subscriber parameters for this SAP.

## non-sub-traffic

**Syntax**        **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*]

**no non-sub-traffic**

**Context**        config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt>single-sub

**Description**    This command configures non-subscriber traffic profiles. It is used in conjunction with the **profiled-traffic-only** command on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues.

The **no** form of the command removes the profiles and disables the feature.

**Parameters**    **sub-profile** *sub-profile-name* — Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

**sla-profile** *sla-profile-name* — Specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

**subscriber** *sub-ident-string* — Specifies an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

For VPRN SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's *sub-ident-string* is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPRN destinations.

If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.

If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

ARP requests are never forwarded back to the same SAP or within the receiving SAP's split horizon group.

## profiled-traffic-only

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] profiled-traffic-only</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt>single-sub                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | This command enables profiled traffic only for this SAP. The profiled traffic refers to single subscriber traffic on a dedicated SAP (in the VLAN-per-subscriber model). When enabled, subscriber queues are instantiated through the QOS policy defined in the sla-profile and the associated SAP queues are deleted. This can increase subscriber scaling by reducing the number of queues instantiated per subscriber (in the VLAN-per-subscriber model). In order for this to be achieved, any configured multi-sub-sap limit must be removed (leaving the default of 1). |

The **no** form of the command disables the command.

## accounting-policy

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>accounting-policy acct-policy-id</b><br><b>no accounting-policy</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>service>ies>if>sap<br>config>service>ies>sub-if>grp-if>sap                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | This command creates the accounting policy context that can be applied to a SAP.<br><br>An accounting policy must be defined before it can be associated with a SAP. If the <i>policy-id</i> does not exist, an error message is generated.<br><br>A maximum of one accounting policy can be associated with a SAP at one time. Accounting policies are configured in the <b>config&gt;log</b> context.<br><br>The <b>no</b> form of this command removes the accounting policy association from the SAP, and the accounting policy reverts to the default. |
| <b>Default</b>     | Default accounting policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

---

|                   |                                                                                                                                 |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>acct-policy-id</i> — The accounting <i>policy-id</i> as configured in the <b>config&gt;log&gt;accounting-policy</b> context. |
| <b>Values</b>     | 1 to 99                                                                                                                         |

## collect-stats

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] collect-stats</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>if>sap<br>config>service>ies>sub-if>grp-if>sap                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>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.</p> <p>When the <b>no collect-stats</b> command is issued the statistics are still accumulated by the IOMCFM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent <b>collect-stats</b> command is issued then the counters written to the billing file include all the traffic while the <b>no collect-stats</b> command was in effect.</p> |
| <b>Default</b>     | no collect-stats                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## bandwidth

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>bandwidth</b> <i>bandwidth</i><br><b>no bandwidth</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>if>sap                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command specifies the admin bandwidth assigned to SAPs, ports and LAGs which is used by SAP bandwidth CAC.</p> <p>SAP: Attempts to increase the SAP admin bandwidth will fail if there is insufficient available admin bandwidth on its port or LAG, otherwise the port or LAG available admin bandwidth will be reduced by the incremental SAP admin bandwidth. Reducing the SAP admin bandwidth will increase the available admin bandwidth on its port or LAG. This is not supported for PW-SAPs, Ethernet tunnels or subscriber group interface SAPs.</p> <p>The <b>no</b> version of the command reverts to the default value.</p> |
| <b>Default</b>     | no bandwidth                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b>  | <i>bandwidth</i> — Specifies the admin bandwidth assigned to the SAP, port or LAG, in kb/s.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Values</b>      | 1 to 3200000000                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## calling-station-id

|                    |                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>calling-station-id</b> <i>calling-station-id</i><br><b>no calling-station-id</b>                                                                                                                                                                           |
| <b>Context</b>     | config>service>ies>if>sap<br>config>service>ies>sub-if>grp-if>sap                                                                                                                                                                                             |
| <b>Description</b> | This command enables the inclusion of the <b>calling-station-id</b> attribute in RADIUS authentication requests and RADIUS accounting messages. The value inserted is set at the SAP level. If no value is set at the SAP level, an empty string is included. |
| <b>Default</b>     | This attribute is not sent by default.                                                                                                                                                                                                                        |

## cpu-protection

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>cpu-protection</b> <i>policy-id</i> {[ <b>mac-monitoring</b> ]   [ <b>eth-cfm-monitoring</b> [ <b>aggregate</b> ] [ <b>car</b> ]]   [ <b>ip-src-monitoring</b> ]}<br><b>no cpu-protection</b>                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>service>ies>if>sap                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command assigns an existing CPU protection policy to the associated service group interface SAP, interface or MSAP policy. The CPU protection policies are configured in the <b>config&gt;sys&gt;security&gt;cpu-protection&gt;policy</b> <i>cpu-protection-policy-id</i> context.</p> <p>If no CPU protection policy is assigned to a service group interface SAP, then a the default policy is used to limit the overall-rate.</p> <p>The <b>no</b> version of this command returns the interface/SAP to the default policies.</p> |
| <b>Default</b>     | <p>cpu-protection 254 (for access interfaces)</p> <p>cpu-protection 255 (for network interfaces)</p> <p>no cpu-protection (for video-interfaces)</p>                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <p><i>policy-id</i> — Specifies an existing CPU protection policy</p> <p><b>Values</b> 1 to 255</p> <p><b>mac-monitoring</b> — When specified, the per MAC rate limiting should be performed, using the per-source-rate from the associated cpu-protection policy</p>                                                                                                                                                                                                                                                                        |

## default-host

|               |                                                                                                                                |
|---------------|--------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>default-host</b> <i>ip-address/mask</i> <b>next-hop</b> <i>next-hop-ip</i><br><b>no default-host</b> <i>ip-address/mask</i> |
|---------------|--------------------------------------------------------------------------------------------------------------------------------|

---

|                    |                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap                                                                                                                                                           |
| <b>Description</b> | This command configures the default-host to be used. More than one default-host can be configured per SAP.<br><br>The <b>no</b> form of the command removes the values from the configuration. |
| <b>Parameters</b>  | <i>ip-address/mask</i> — Assigns an IP address/IP subnet format to the interface<br><b>next-hop</b> <i>next-hop-ip</i> — Assigns the next hop IP address                                       |

## dist-cpu-protection

|                    |                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dist-cpu-protection</b> <i>policy-name</i><br><b>no dist-cpu-protection</b>                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap<br>config>service>ies>if>sap                                                                                                                                                                                                     |
| <b>Description</b> | This command assigns a Distributed CPU Protection (DCP) policy to the SAP. Only a valid DCP policy can be assigned to a SAP or a network interface. This rule does not apply to templates such as an msap-policy.                                                     |
| <b>Default</b>     | If no dist-cpu-protection policy is assigned to an SAP, then the default access DCP policy (default-access-policy) is used.<br><br>If no DCP functionality is required on the SAP, then an empty DCP policy can be created and explicitly assigned to the SAP policy. |
| <b>Parameters</b>  | <i>policy-name</i> — Specifies the name of the DCP policy up to 32 characters in length                                                                                                                                                                               |

## 2.5.2.14 ETH-CFM Service Commands

### eth-cfm

|                    |                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>eth-cfm</b>                                                                                                              |
| <b>Context</b>     | config>service>ies><br>config>service>ies>sub-if>grp-if>sap<br>config>service>ies>if>sap<br>config>service>ies>if>spoke-sdp |
| <b>Description</b> | This command enters the context to configure ETH-CFM parameters.                                                            |



## collect-lmm-fc-stats

|                    |                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>collect-lmm-fc-stats</b>                                                                                                                                                                                                      |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm<br>config>service>ies>if>spoke-sdp>eth-cfm<br>config>service>ies>sub-if>grp-if>sap>eth-cfm                                                                                                     |
| <b>Description</b> | This command enters the context to configure per-forwarding class (FC) LMM information collection.<br><br>This command is mutually exclusive with the <b>collect-lmm-stats</b> command when there is entity resource contention. |

## collect-lmm-stats

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>collect-lmm-stats</b><br><b>no collect-lmm-stats</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm<br>config>service>ies>if>spoke-sdp>eth-cfm<br>config>service>ies>sub-if>grp-if>sap>eth-cfm                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | This command enables the collection of statistics on the SAP or MPLS SDP binding on which the ETH- LMM test is configured. The collection of LMM statistics must be enabled if a MEP is launching or responding to ETH-LMM packets. If LMM statistics collection is not enabled, the counters in the LMM and LMR PDU do not represent accurate measurements and all measurements should be ignored. The <b>show sap-using eth-cfm collect-lmm-stats</b> command and the <b>show sdp-using eth-cfm collect-lmm-stats</b> command can be used to display which entities are collecting stats.<br><br>The <b>no</b> form of the command disables and deletes the counters for this SAP or MPLS SDP binding. |
| <b>Default</b>     | no collect-lmm-stats                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## fc

|                    |                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fc fc-name [fc-name]</b><br><b>no fc</b>                                                                                                                                                           |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>collect-lmm-fc-stats<br>config>service>ies>if>spoke-sdp>eth-cfm>collect-lmm-fc-stats<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>collect-lmm-fc-stats           |
| <b>Description</b> | This command creates individual counters for the specified FCs without regard for profile. All countable packets that match a configured FC, regardless of profile, will be included in this counter. |

A differential is performed when this command is re-entered. Omitted FCs will stop counting, newly added FCs will start counting, and unchanged FCs will continue to count.

An FC that is specified as part of this command for this specific context cannot be specified as a profile-aware FC using the **fc-in-profile** command under the same context.

The **no** form of the command removes all previously defined FCs and stops counting for those FCs.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b> | <p><i>fc-name</i> — Specifies the name of the FC for which to create an individual profile-unaware counter. Up to eight FCs may be specified. In order for the counter to be used, the <b>config&gt;oam-pm&gt;session&gt;ethernet&gt;priority</b> command must be configured with a numerical value representing the FC name (7 = NC, 6 = H1, 5 = EF, 4 = H2, 3 = L1, 2 = AF, 1 = L2, 0 = BE), and the <b>config&gt;oam-pm&gt;session&gt;ethernet&gt;lmm&gt;enable-fc-collection</b> command must be enabled.</p> <p><b>Values</b>      nc, h1, ef, h2, l1, af, l2, be</p> |

## fc-in-profile

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fc-in-profile</b> <i>fc-name</i> [ <i>fc-name</i> ]<br><b>no fc-in-profile</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>collect-lmm-fc-stats<br>config>service>ies>if>spoke-sdp>eth-cfm>collect-lmm-fc-stats<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>collect-lmm-fc-stats                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command creates individual counters for the specified FCs with regard for profile. All countable packets that match a configured FC and are deemed to be in profile will be included in this counter.</p> <p>A differential is performed when this command is re-entered. Omitted FCs will stop counting, newly added FCs will start counting, and unchanged FCs will continue to count.</p> <p>An FC that is specified as part of this command for this specific context cannot be specified as a profile-unaware FC using the <b>fc</b> command under the same context.</p> <p>The <b>no</b> form of the command removes all previously defined FCs and stops counting for those FCs.</p> |
| <b>Default</b>     | no fc-in-profile                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

**Parameters** *fc-name* — Specifies the name of the FC for which to create an individual profile-aware counter. Up to eight FCs may be specified. In order for the counter to be used, the **config>oam-pm>session>ethernet>priority** command must be configured with a numerical value representing the FC name (7 = NC, 6 = H1, 5 = EF, 4 = H2, 3 = L1, 2 = AF, 1 = L2, 0 = BE), and the **config>oam-pm>session>ethernet>imm>enable-fc-collection** command must be enabled.

**Values** nc, h1, ef, h2, l1, af, l2, be

## 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>ies>if>sap>eth-cfm  
config>service>ies>if>spoke-sdp>eth-cfm  
config>service>ies>sub-if>grp-if>sap>eth-cfm

**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** — The direction in which the maintenance association (MEP) faces on the bridge port. Direction UP is not applicable to IES MEPs.  
**down** — Sends ETH-CFM messages away from the MAC relay entity.  
**up** — Sends ETH-CFM messages towards the MAC relay entity.

## ais-enable

**Syntax** [**no**] **ais-enable**

**Context** config>service>ies>if>spoke-sdp>eth-cfm

**Description** This command configures the reception of Alarm Indication Signal (AIS) message.

## interface-support-enable

**Syntax** [**no**] **interface-support-enable**

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>ies>sap>eth-cfm>mep>ais-enable<br>config>service>ies>spoke-sdp>eth-cfm>mep>ais-enable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | This command enables the AIS function to consider the operational state of the entity on which it is configured. With this command, ETH-AIS on DOWN MEPs will be triggered and cleared based on the operational status of the entity on which it is configured. If CCM is also enabled then transmission of the AIS PDU will be based on either the non-operational state of the entity or on any CCM defect condition. AIS generation will cease if BOTH operational state is UP and CCM has no defect conditions. If the MEP is not CCM enabled then the operational state of the entity is the only consideration assuming this command is present for the MEP. |
| <b>Default</b>     | no interface-support-enable (AIS will not be generated or stopped based on the state of the entity on which the DOWN MEP is configured).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## ccm-enable

|                    |                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ccm-enable</b>                                                                                                                   |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep<br>config>service>ies>if>spoke-sdp>eth-cfm>mep<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep |
| <b>Description</b> | This command enables the generation of CCM messages.<br><br>The <b>no</b> form of the command disables the generation of CCM messages.   |

## ccm-ltm-priority

|                    |                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ccm-ltm-priority</b> <i>priority</i><br><b>no ccm-ltm-priority</b>                                                                                                           |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep<br>config>service>ies>if>spoke-sdp>eth-cfm>mep<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep                                        |
| <b>Description</b> | This command specifies the priority value for CCMs and LTMs transmitted by the MEP.<br><br>The <b>no</b> form of the command removes the priority value from the configuration. |
| <b>Default</b>     | The highest priority on the bridge-port.                                                                                                                                        |
| <b>Parameters</b>  | <i>priority</i> — Specifies the priority of CCM and LTM messages.<br><br><b>Values</b> 0 to 7                                                                                   |

## ccm-padding-size

|                    |                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ccm-padding-size</b> <i>ccm-padding</i><br><b>no ccm-padding-size</b> <i>ccm-padding</i>                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep<br>config>service>ies>if>spoke-sdp>eth-cfm>mep                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | Set the byte size of the optional Data TLV to be included in the ETH-CC PDU. This will increase the size of the ETH-CC PDU by the configured value. The base size of the ETH-CC PDU, including the Interface Status TLV and Port Status TLV, is 83 bytes not including the Layer Two encapsulation. CCM padding is not supported when the CCM-Interval is less than one second. |
| <b>Default</b>     | ccm-padding-size                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <i>ccm-padding</i> — Specifies the byte size of the Optional Data TLV.<br><b>Values</b> 3 to 1500                                                                                                                                                                                                                                                                               |

## eth-test-enable

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] eth-test-enable</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep<br>config>service>ies>if>spoke-sdp>eth-cfm>mep<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | For ETH-test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands:<br><br>oam eth-cfm eth-test <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [priority <i>priority</i> ] [data-length <i>data-length</i> ]<br><br>A check is done for both the provisioning and test to ensure the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the CLI and SNMP will indicate the problem. |

## bit-error-threshold

|                    |                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>bit-error-threshold</b> <i>bit-errors</i>                                                                         |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep>eth-test-enable<br>config>service>ies>if>spoke-sdp>eth-cfm>mep>eth-test-enable |
| <b>Description</b> | This command specifies the lowest priority defect that is allowed to generate a fault alarm.                         |
| <b>Default</b>     | bit-error-threshold 1                                                                                                |

---

|                   |                                                           |
|-------------------|-----------------------------------------------------------|
| <b>Parameters</b> | <i>bit-errors</i> — Specifies the lowest priority defect. |
| <b>Values</b>     | 0 to 11840                                                |

## test-pattern

|                    |                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>test-pattern {all-zeros   all-ones} [crc-enable]</b><br><b>no test-pattern</b>                                                                                                                                   |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep>eth-test-enable<br>config>service>ies>if>spoke-sdp>eth-cfm>mep>eth-test-enable<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep>eth-test-enable                            |
| <b>Description</b> | This command configures the test pattern for eth-test frames.<br><br>The <b>no</b> form of the command removes the values from the configuration.                                                                   |
| <b>Parameters</b>  | <b>all-zeros</b> — Specifies to use all zeros in the test pattern.<br><b>all-ones</b> — Specifies to use all ones in the test pattern.<br><b>crc-enable</b> — Generates a CRC checksum.<br><b>Default</b> all-zeros |

## fault-propagation-enable

|                    |                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fault-propagation-enable {use-if-tlv   suspend-ccm}</b><br><b>no fault-propagation-enable</b>                                         |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep<br>config>service>ies>if>spoke-sdp>eth-cfm>mep<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep |
| <b>Description</b> | This command configures the fault propagation for the MEP.                                                                               |
| <b>Parameters</b>  | <b>use-if-tlv</b> — Specifies to use the interface TLV.<br><b>suspend-ccm</b> — Specifies to suspend the continuity check messages.      |

## grace

|                |                                                                                                                                          |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>grace</b>                                                                                                                             |
| <b>Context</b> | config>service>ies>if>sap>eth-cfm>mep<br>config>service>ies>if>spoke-sdp>eth-cfm>mep<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep |

---

|                    |                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command enters the context to configure Nokia ETH-CFM Grace and ITU-T Y.1731 ETH-ED expected defect functional parameters. |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------|

## eth-ed

|                    |                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>eth-ed</b>                                                                                                                                              |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep>grace<br>config>service>ies>if>spoke-sdp>eth-cfm>mep>grace<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep>grace |
| <b>Description</b> | This command enters the context to configure ITU-T Y.1731 ETH-ED expected defect functional parameters.                                                    |

## max-rx-defect-window

|                    |                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>max-rx-defect-window</b> <i>seconds</i><br><b>no max-rx-defect-window</b>                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep>grace>eth-ed<br>config>service>ies>if>spoke-sdp>eth-cfm>mep>grace>eth-ed<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep>grace>eth-ed                                                                                                                                             |
| <b>Description</b> | This command limits the duration of the received ETH-ED expected defect window to the lower value of either the received value from the peer or this parameter.<br><br>The <b>no</b> form of the command removes the limitation, and any valid defect window value received from a peer MEP in the ETH-ED PDU will be used. |
| <b>Default</b>     | no max-rx-defect-window                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <i>seconds</i> — Specifies the duration, in seconds, of the maximum expected defect window<br><br><b>Values</b> 1 to 86400                                                                                                                                                                                                  |

## priority

|                    |                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>priority</b> <i>priority</i><br><b>no priority</b>                                                                                                                           |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep>grace>eth-ed<br>config>service>ies>if>spoke-sdp>eth-cfm>mep>grace>eth-ed<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep>grace>eth-ed |
| <b>Description</b> | This command sets the priority bits and determines the forwarding class based on the mapping of priority to FC.                                                                 |

The **no** form of the command disables the local priority configuration and sets the priority to the **ccm-ltm-priority** associated with this MEP.

|                   |                                               |
|-------------------|-----------------------------------------------|
| <b>Default</b>    | no priority                                   |
| <b>Parameters</b> | <i>priority</i> — Specifies the priority bit. |
| <b>Values</b>     | 0 to 7                                        |

## rx-eth-ed

|                    |                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] rx-eth-ed</b>                                                                                                                                                                                   |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep>grace>eth-ed<br>config>service>ies>if>spoke-sdp>eth-cfm>mep>grace>eth-ed<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep>grace>eth-ed                         |
| <b>Description</b> | This command enables the reception and processing of the ITU-T Y.1731 ETH-ED PDU on the MEP.<br><br>The <b>no</b> form of the command disables the reception of the ITU-T Y.1731 ETH-ED PDU on the MEP. |
| <b>Default</b>     | rx-eth-ed                                                                                                                                                                                               |

## tx-eth-ed

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] tx-eth-ed</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep>grace>eth-ed<br>config>service>ies>if>spoke-sdp>eth-cfm>mep>grace>eth-ed<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep>grace>eth-ed                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | This command enables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP when a system soft reset notification is received for one or more cards.<br><br>The <b>config&gt;eth-cfm&gt;system&gt;grace-tx-enable</b> command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted.<br><br>The <b>no</b> form of the command disables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP. |
| <b>Default</b>     | no tx-eth-ed                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



## eth-vsm-grace

|                    |                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>eth-vsm-grace</b>                                                                                                                                       |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep>grace<br>config>service>ies>if>spoke-sdp>eth-cfm>mep>grace<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep>grace |
| <b>Description</b> | This command enters the context to configure Nokia ETH-CFM Grace functional parameters.                                                                    |

## rx-eth-vsm-grace

|                    |                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] rx-eth-vsm-grace</b>                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep>grace>eth-vsm-grace<br>config>service>ies>if>spoke-sdp>eth-cfm>mep>grace>eth-vsm-grace<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep>grace>eth-vsm-grace                                                                                                                                                            |
| <b>Description</b> | <p>This command enables the reception and processing of the Nokia ETH-CFM Grace PDU on the MEP.</p> <p>The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect.</p> <p>The <b>no</b> form of the command disables the reception of the Nokia ETH-CFM Grace PDU on the MEP.</p> |
| <b>Default</b>     | rx-eth-vsm-grace                                                                                                                                                                                                                                                                                                                                                |

## tx-eth-vsm-grace

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] tx-eth-vsm-grace</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep>grace>eth-vsm-grace<br>config>service>ies>if>spoke-sdp>eth-cfm>mep>grace>eth-vsm-grace<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep>grace>eth-vsm-grace                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command enables the transmission of the Nokia ETH-CFM Grace PDU from the MEP when a system soft reset notification is received for one or more cards.</p> <p>The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect.</p> <p>The <b>config&gt;eth-cfm&gt;system&gt;grace-tx-enable</b> command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted.</p> |

The **no** form of the command disables the transmission of the Nokia ETH-CFM Grace PDU from the MEP.

**Default** tx-eth-vsm-grace

## low-priority-defect

|                    |                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>low-priority-defect</b> {allDef   macRemErrXcon   remErrXcon   errXcon   xcon   noXcon}                                                  |
| <b>Context</b>     | config>service>ies>if>sap>eth-cfm>mep<br>config>service>ies>if>spoke-sdp>eth-cfm>mep<br>config>service>ies>sub-if>grp-if>sap>eth-cfm>mep    |
| <b>Description</b> | This command specifies the lowest priority defect that is allowed to generate a fault alarm.                                                |
| <b>Default</b>     | low-priority-defect macRemErrXcon                                                                                                           |
| <b>Parameters</b>  | <b>low-priority-defect</b> — The following values are used to specify the lowest priority defect that is allowed to generate a fault alarm. |
| <b>Values</b>      |                                                                                                                                             |
| allDef             | DefRDICCM, DefMACstatus, DefRemoteCCM, efErrorCCM, 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                                                                                              |

## squelch-ingress-levels

|                |                                                                                                                              |
|----------------|------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>squelch-ingress-levels</b> [ <i>md-level</i> [ <i>md-level</i> ]]<br><b>no squelch-ingress-levels</b>                     |
| <b>Context</b> | config>service>ies>if>sap>eth-cfm<br>config>service>ies>if>spoke-sdp>eth-cfm<br>config>service>ies>sub-if>grp-if>sap>eth-cfm |

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command defines the levels of the ETH-CFM PDUs that will silently be discarded on ingress into the SAP or SDP Binding from the wire. All ETH-CFM PDUs inbound to the SAP or SDP binding will be dropped that match the configured levels without regard for any other ETH-CFM criteria. No statistical information or drop count will be available for any ETH-PDU that is silently discarded by this option. The operator must configure a complete contiguous list of md-levels up to the highest level that will be dropped. The command must be retyped in complete form to modify a previous configuration, if the operator does not want to delete it first.</p> <p>The <b>no</b> form of the command removes the silent discarding of previously matching ETH-CFM PDUs.</p> |
| <b>Default</b>     | no squelch-ingress-levels                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <i>md-level</i> — Identifies the level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|                    | <b>Values</b> 0 to 7                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## tunnel-fault

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>tunnel-fault {accept   ignore}</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>service>ies>eth-cfm<br>config>service>ies>if>sap>eth-cfm<br>config>service>ies>sub-if>grp-if>sap>eth-cfm                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>Allows the individual service SAPs to react to changes in the tunnel MEP state. When tunnel-fault accept is configured at the service level, the SAP will react according to the service type, Epipe will set the operational flag and VPLS, IES and VPRN SAP operational state will become down on failure or up on clear. This command triggers the OAM mapping functions to mate SAPs and bindings in an Epipe service as well as setting the operational flag. If AIS generation is the requirement for the Epipe services this command is not required. See the <b>ais-enable</b> command in the <b>epipe&gt;sap&gt;eth-cfm</b> context for more information. This works in conjunction with the tunnel-fault accept on the individual SAPs. Both must be set to accept to react to the tunnel MEP state. By default the service level command is <b>ignore</b> and the sap level command is <b>accept</b>. This means simply changing the service level command to <b>accept</b> enables the feature for all SAPs. This is not required for Epipe services that only wish to generate AIS on failure.</p> |
| <b>Default</b>     | tunnel-fault ignore (Service Level)<br><br>tunnel-fault accept (SAP Level for Epipe and VPLS)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <i>accept</i> — Share fate with the facility tunnel MEP.<br><i>ignore</i> — Do not share fate with the facility tunnel MEP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## one-way-delay-threshold

|                    |                                                                              |
|--------------------|------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>one-way-delay-threshold</b> <i>time</i>                                   |
| <b>Context</b>     | config>service>ies>if>sap>mep<br>config>service>ies>if>spoke-sdp>eth-cfm>mep |
| <b>Description</b> | This command enables one way delay threshold time limit.                     |
| <b>Default</b>     | 3 seconds                                                                    |
| <b>Parameters</b>  | <i>priority</i> — Specifies the value for the threshold.                     |
|                    | <b>Values</b> 0 to 600                                                       |

### 2.5.2.15 IES Filter and QoS Policy Commands

## filter

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>filter ip</b> <i>ip-filter-id</i><br><b>filter ipv6</b> <i>ipv6-filter-id</i><br><b>no filter</b> [ <b>ip</b> <i>ip-filter-id</i> ] [ <b>ipv6</b> <i>ipv6-filter-id</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>service>ies>if>sap>egress<br>config>service>ies>if>sap>ingress<br>config>service>ies>red-if>egress<br>config>service>ies>red-if>ingress<br>config>service>ies>red-if>egress<br>config>service>ies>red-if>ingress<br>config>service>ies>sub-if>grp-if>sap>egress<br>config>service>ies>sub-if>grp-if>sap>ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command associates a filter policy with an ingress or egress Service Access Point (SAP). Filter policies control the forwarding and dropping of packets based on the matching criteria.</p> <p>The <b>filter</b> command is used to associate a filter policy with a specified <i>ip-filter-id</i> or <i>ipv6-filter-id</i> (7750 SR) with an ingress or egress SAP. The filter policy must already be defined before the <b>filter</b> command is executed. If the filter policy does not exist, the operation fails and an error message returned.</p> <p>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.</p> |

The **no** form of this command removes any configured filter ID association with the SAP. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

**Special Cases** **IES** — Only IP filters are supported on an IES IP interface, and the filters only apply to routed traffic.

**Parameters** **ip** — indicates the filter policy is an IP filter.  
*ip-filter-id* — Specifies the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy in the **config>filter>ip-filter** context.

## filter

**Syntax** **filter ip** *ip-filter-id*  
**filter ipv6** *ipv6-filter-id*  
**no filter**

**Context** config>service>ies>if>spoke-sdp>egress  
config>service>ies>if>spoke-sdp>ingress

**Description** This command associates an IP filter policy filter policy with an ingress or egress spoke SDP.

Filter policies control the forwarding and dropping of packets based on matching criteria.

MAC filters are only allowed on Epipe and Virtual Private LAN Service (VPLS) SAPs.

The **filter** command is used to associate a filter policy with a specified *ip-filter-id* with an ingress or egress spoke SDP. The *ip-filter-id* must already be defined in the **config>filter** context before the **filter** command is executed. If the filter policy does not exist, the operation fails and an error message returned.

In general, filters applied to SAPs or spoke SDPs (ingress or egress) apply to all packets on the SAP or spoke SDPs. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

**Special Cases** **IES** — Only IP filters are supported on IES IP interfaces, and the filters only apply to routed traffic.

**Parameters** **ip** — Keyword indicating the filter policy is an IP filter.

---

*ip-filter-id* — The filter name acts as 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. The filter ID must already exist within the created IP filters.

## egress

|                    |                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>egress</b>                                                                                                                                                                                                                               |
| <b>Context</b>     | config>service>ies>if>sap<br>config>service>ies>sub-if>grp-if>sap                                                                                                                                                                           |
| <b>Description</b> | This command enters the context to apply egress policies.<br><br>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

|                    |                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ingress</b>                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>service>ies>if>sap<br>config>service>ies>sub-if>grp-if>sap                                                                                                                                                                                |
| <b>Description</b> | This command enters the context to apply ingress policies.<br><br>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. |

## hsmda-queue-override

|                    |                                                                       |
|--------------------|-----------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] hsmda-queue-override</b>                                      |
| <b>Context</b>     | config>service>ies>if>sap>egress<br>config>service>vprn>if>sap>egress |
| <b>Description</b> | This command configures HSM DA egress and ingress queue overrides.    |

## packet-byte-offset

|                |                                                                                                              |
|----------------|--------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>packet-byte-offset {add <i>add-bytes</i>   subtract <i>sub-bytes</i>}</b><br><b>no packet-byte-offset</b> |
| <b>Context</b> | config>service>ies>if>sap>egress>hsmda-queue-over                                                            |

**Description** This command adds or subtracts the specified number of bytes to the accounting function for each packet handled by the HSMDA queue. Normally, the accounting and leaky bucket functions are based on the Ethernet DLC header, payload and the 4-byte CRC (everything except the preamble and inter-frame gap). For example, this command can be used to add the frame encapsulation overhead (20 bytes) to the queues accounting functions.

The accounting functions affected include:

- Offered High Priority / In-Profile Octet Counter
- Offered Low Priority / Out-of-Profile Octet Counter
- Discarded High Priority / In-Profile Octet Counter
- Discarded Low Priority / Out-of-Profile Octet Counter
- Forwarded In-Profile Octet Counter
- Forwarded Out-of-Profile Octet Counter
- Peak Information Rate (PIR) Leaky Bucket Updates
- Committed Information Rate (CIR) Leaky Bucket Updates
- Queue Group Aggregate Rate Limit Leaky Bucket Updates

The secondary shaper leaky bucket, scheduler priority level leaky bucket and the port maximum rate updates are not affected by the configured packet-byte-offset. Each of these accounting functions are frame based and always include the preamble, DLC header, payload and the CRC regardless of the configured byte offset.

The packet-byte-offset command accepts either add or subtract as valid keywords which define whether bytes are being added or removed from each packet traversing the queue. Up to 20 bytes may be added to the packet and up to 43 bytes may be removed from the packet. An example use case for subtracting bytes from each packet is an IP based accounting function. Given a Dot1Q encapsulation, the command packet-byte-offset subtract 14 would remove the DLC header and the Dot1Q header from the size of each packet for accounting functions only. The 14 bytes are not actually removed from the packet, only the accounting size of the packet is affected.

As mentioned above, the variable accounting size offered by the packet-byte-offset command is targeted at the queue and queue group level. When the queue group represents the last-mile bandwidth constraints for a subscriber, the offset allows the HSMDA queue group to provide an accurate accounting to prevent overrun and underrun conditions for the subscriber. The accounting size of the packet is ignored by the secondary shapers, the scheduling priority level shapers and the scheduler maximum rate. The actual on-the-wire frame size is used for these functions to allow an accurate representation of the behavior of the subscriber's packets on an Ethernet aggregation network.

The packet-byte-offset value can be overridden for the HSMDA queue at the SAP or subscriber profile level.

The **no** form of the command removes any accounting size changes to packets handled by the queue. The command does not affect overrides that may exist on SAPs or subscriber profiles associated with the queue.

---

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <p><b>add</b> <i>add-bytes</i> — The <b>add</b> keyword is mutually exclusive with the subtract keyword. Either the add or subtract keyword must be specified. The add keyword is used to indicate that the following byte value should be added to the packet for queue and queue group level accounting functions. The corresponding byte value must be specified when executing the packet-byte-offset command.</p> <p><b>Values</b>      0 to 31</p> <p><b>subtract</b> <i>sub-bytes</i> — The <b>subtract</b> keyword is mutually exclusive with the add keyword. Either the add or subtract keyword must be specified. The subtract keyword is used to indicate that the following byte value should be subtracted from the packet for queue and queue group level accounting functions. The corresponding byte value must be specified when executing the packet-byte-offset command.</p> <p><b>Values</b>      1 to 64</p> |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## queue

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>queue</b> <i>queue-id</i> [<b>create</b>]</p> <p><b>no queue</b> <i>queue-id</i></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>service>ies>if>sap>egress>hsmdda-queue-over                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command, within the QoS policy hsmdda-queue context, is a container for the configuration parameters controlling the behavior of an HSMDDA queue. Unlike the standard QoS policy queue command, this command is not used to actually create or dynamically assign the queue to the object which the policy is applied. The queue identified by queue-id always exists on the SAP or subscriber context whether the command is executed or not. In the case of HSMDDA SAPs and subscribers, all eight queues exist at the moment the system allocates an HSMDDA queue group to the object (both ingress and egress).</p> |

### Best-Effort, Expedited and Auto-Expedite Queue Behavior Based on Queue-ID

With standard service queues, the scheduling behavior relative to other queues is based on two items, the queues Best-Effort or Expedited nature and the dynamic rate of the queue relative to the defined CIR. HSMDDA queues are handled differently. The create time auto-expedite and explicit expedite and best-effort qualifiers have been eliminated and instead the scheduling behavior is based solely on the queues identifier. Queues with a queue-id equal to 1 are placed in scheduling class 1. Queues with queue-id 2 are placed in scheduling class 2. And so on up to scheduling class 8. Each scheduling class is either mapped directly to a strict scheduling priority level based on the class ID, or the class may be placed into a weighted scheduling class group providing byte fair weighted round robin scheduling between the members of the group. Two weighted groups are supported and each may contain up to three consecutive scheduling classes. The weighed group assumes its highest member classes inherent strict scheduling level for scheduling purposes. Strict priority level 8 has the highest priority; strict level 1 has the lowest priority. When grouping of scheduling classes is defined, some of the strict levels will not be in use.

### Single Type of HSMDDA Queues



Another difference between HSMDB queues and standard service queues is the lack of Multipoint queues. At ingress, an HSMDB SAP or subscriber does not require Multipoint queues since all forwarding types (broadcast, multicast, unicast and unknown) forward to a single destination in the ingress forwarding plane on the IOM. Instead of a possible eight queues per forwarding type (for a total of up to 32) within the SAP ingress QoS policy, the hsmdb-queues node supports a maximum of eight queues.

#### Every HSMDB Queue Supports Profile Mode Implicitly

Unlike standard service queues, the HSMDB queues do not need to be placed into the special mode profile at create time in order to support ingress color aware policing. Each queue may handle in-profile, out-of-profile and profile undefined packets simultaneously. As with standard queues, the explicit profile of a packet is dependent on the ingress sub-forwarding class to which the packet is mapped.

The **no** form of the command restores the defined queue-id to its default parameters. All HSMDB queues having the queue-id and associated with the QoS policy are re-initialized to default parameters.

**Parameters** *queue-id* — Specifies the HSMDB queue to use for packets in this forwarding class. This mapping is used when the SAP is on a HSMDB MDA.

**Values** 1 to 8

## rate

**Syntax** **rate** *pir-rate*  
**no rate**

**Context** config>service>ies>if>sap>egress>hsmdb-queue-over>queue

**Description** This command specifies the administrative PIR by the user.

**Parameters** *pir-rate* — Configures the administrative PIR specified by the user.

**Values** 1 to 100000000, **max**

## slope-policy

**Syntax** **slope-policy** *hsmdb-slope-policy-name*  
**no slope-policy**

**Context** config>service>ies>if>sap>egress>hsmdb-queue-over

**Description** This command assigns an HSMDBA slope policy to the SAP. The policy may be assigned to an ingress or egress HSMDBA queue. The policy contains the Maximum Buffer Size (MBS) that will be applied to the queue and the high and low priority RED slope definitions. The function of the MBS and RED slopes is to provide congestion control for an HSMDBA queue. The MBS parameter defines the maximum depth a queue may reach when accepting packets. The low and high priority RED slopes provides for random early detection of congestion and slope based discards based on queue depth.

An HSMDBA slope policy can be applied to queues defined in the SAP ingress and SAP egress QoS policy HSMDBA queues context. Once an HSMDBA slope policy is applied to a SAP QoS policy queue, it cannot be deleted. Any edits to the policy are updated to all HSMDBA queues indirectly associated with the policy.

#### Default HSMDBA Slope Policy

An HSMDBA slope policy named **default** always exists on the system and does not need to be created. The default policy is automatically applied to all HSMDBA queues unless another HSMDBA slope policy is specified for the queue. The default policy cannot be modified or deleted. Attempting to execute the **no hsmdb-slope-policy default** command results in an error.

The **no** form of the command removes the specified HSMDBA slope policy from the configuration. If the HSMDBA slope policy is currently associated with an HSMDBA queue, the command fails.

**Parameters** *hsmdb-slope-policy-name* — Specifies a HSMDBA slope policy up to 32 characters in length. The HSMDBA slope policy must exist prior to applying the policy name to an HSMDBA queue.

## wrr-weight

**Syntax** **wrr-weight value**  
**no wrr-weight**

**Context** config>service>ies>if>sap>egress>hsmdb-queue-override>queue

**Description** This command assigns the weight value to the HSMDBA queue.

The **no** form of the command returns the weight value for the queue to the default value.

**Parameters** *percentage* — Specifies the weight for the HSMDBA queue.

**Values** 1 to 32

## wrr-policy

**Syntax** **wrr-policy hsmdb-wrr-policy-name**

---

|                    |                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>ies>if>sap>egress>hsmdda-queue-override                                                            |
| <b>Description</b> | This command associates an existing HSMDDA weighted-round-robin (WRR) scheduling loop policy to the HSMDDA queue. |
| <b>Parameters</b>  | <i>hsmdda-wrr-policy-name</i> — Specifies the existing HSMDDA WRR policy name to associate to the queue.          |

## secondary-shaper

|                    |                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>secondary-shaper</b> <i>secondary-shaper-name</i><br><b>no secondary-shaper</b>              |
| <b>Context</b>     | config>service>ies>if>sap>egress>hsmdda-queue-override                                          |
| <b>Description</b> | This command configures an HSMDDA egress secondary shaper.                                      |
| <b>Parameters</b>  | <i>secondary-shaper-name</i> — Specifies a secondary shaper name up to 32 characters in length. |

## policer-override

|                    |                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] policer-override</b>                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>service>ies>if>sap>egress<br>config>service>ies>if>sap>ingress                                                                                                                                                                                                                                  |
| <b>Description</b> | This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to one or more policers created on the SAP through the sap-ingress or sap-egress QoS policies.<br><br>The <b>no</b> form of the command is used to remove any existing policer overrides. |
| <b>Default</b>     | no policer-override                                                                                                                                                                                                                                                                                    |

## policer

|                    |                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>policer</b> <i>policer-id</i> <b>[create]</b><br><b>no policer</b> <i>policer-id</i>                                                                                                               |
| <b>Context</b>     | config>service>ies>if>sap>egress>policer-override<br>config>service>ies>if>sap>ingress>policer-override                                                                                               |
| <b>Description</b> | This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to a specific policer created on the SAP through a sap-ingress or sap-egress QoS policy. |

The **no** form of the command is used to remove any existing overrides for the specified *policer-id*.

**Parameters** *policer-id* — This parameter is required when executing the *policer* command within the *policer-override* context. The specified *policer-id* must exist within the *sap-ingress* or *sap-egress* QoS policy applied to the SAP. If the *policer* is not currently used by any forwarding class or forwarding type mappings, the *policer* will not actually exist on the SAP. This does not preclude creating an override context for the *policer-id*.

**create** — The *create* keyword is required when a *policer 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 configuration, the **create** keyword is not required.

## match-qinq-dot1p

**Syntax** **match-qinq-dot1p** {**top** | **bottom**}  
**no match-qinq-dot1p**

**Context** config>service>ies>if>sap>ingress  
config>service>ies>sub-if>grp-if>sap>ingress

**Description** This command specifies which Dot1Q tag position Dot1P bits in a QinQ encapsulated packet should be used to evaluate Dot1P QoS classification.

The **match-qinq-dot1p** command allows the *top* or *bottom* PBits to be used when evaluating the applied *sap-ingress* QoS policy's Dot1P entries. The **top** and **bottom** keywords specify which position should be evaluated for QinQ encapsulated packets.

The **no** form of the command restores the default *dot1p* evaluation behavior for the SAP.

By default, the bottom most service delineating Dot1Q tags Dot1P bits are used. [Table 8](#) defines the default behavior for Dot1P evaluation when the **match-qinq-dot1p** command is not executed.

**Table 8** Default QinQ and TopQ SAP Dot1P Evaluation

| Port / SAP Type | Existing Packet Tags | PBits Used for Match |
|-----------------|----------------------|----------------------|
| null            | none                 | none                 |
| null            | Dot1P (VLAN-ID 0)    | Dot1P PBits          |
| null            | Dot1Q                | Dot1Q PBits          |
| null            | TopQ BottomQ         | TopQ PBits           |
| null            | TopQ (No BottomQ)    | TopQ PBits           |
| Dot1Q           | none (Default SAP)   | none                 |

**Table 8 Default QinQ and TopQ SAP Dot1P Evaluation (Continued)**

| Port / SAP Type | Existing Packet Tags          | PBits Used for Match |
|-----------------|-------------------------------|----------------------|
| Dot1Q           | Dot1P (Default SAP VLAN-ID 0) | Dot1P PBits          |
| Dot1Q           | Dot1Q                         | Dot1Q PBits          |
| QinQ / TopQ     | TopQ                          | TopQ PBits           |
| QinQ / TopQ     | TopQ BottomQ                  | TopQ PBits           |
| QinQ / QinQ     | TopQ BottomQ                  | BottomQ PBits        |

**Default** no match-qinq-dot1p — no filtering based on p-bits

top or bottom must be specified to override the default QinQ dot1p behavior.

**Parameters** **top** — The top parameter is mutually exclusive to the bottom parameter. When the top parameter is specified, the top most PBits are used (if existing) to match any dot1p dot1p-value entries. [Table 9](#) defines the dot1p evaluation behavior when the top parameter is specified.

**Table 9 Top Position QinQ and TopQ SAP Dot1P Evaluation**

| Port / SAP Type | Existing Packet Tags          | PBits Used for Match |
|-----------------|-------------------------------|----------------------|
| null            | none                          | none                 |
| null            | Dot1P (VLAN-ID 0)             | Dot1P PBits          |
| null            | Dot1Q                         | Dot1Q PBits          |
| null            | TopQ BottomQ                  | TopQ PBits           |
| null            | TopQ (No BottomQ)             | TopQ PBits           |
| Dot1Q           | none (Default SAP)            | none                 |
| Dot1Q           | Dot1P (Default SAP VLAN-ID 0) | Dot1P PBits          |
| Dot1Q           | Dot1Q                         | Dot1Q PBits          |
| QinQ / TopQ     | TopQ                          | TopQ PBits           |
| QinQ / TopQ     | TopQ BottomQ                  | TopQ PBits           |
| QinQ / QinQ     | TopQ BottomQ                  | TopQ PBits           |

**bottom** — The bottom parameter is mutually exclusive to the top parameter. When the bottom parameter is specified, the bottom most PBits are used (if existing) to match any dot1p dot1p-value entries. [Table 10](#) defines the dot1p evaluation behavior when the bottom parameter is specified.

**Table 10** Bottom Position QinQ and TopQ SAP Dot1P Evaluation

| Port / SAP Type | Existing Packet Tags          | PBits Used for Match |
|-----------------|-------------------------------|----------------------|
| null            | none                          | none                 |
| null            | Dot1P (VLAN-ID 0)             | Dot1P PBits          |
| null            | Dot1Q                         | Dot1Q PBits          |
| null            | TopQ BottomQ                  | BottomQ PBits        |
| null            | TopQ (no BottomQ)             | TopQ PBits           |
| Dot1Q           | none (default SAP)            | none                 |
| Dot1Q           | Dot1P (Default SAP VLAN-ID 0) | Dot1P PBits          |
| Dot1Q           | Dot1Q                         | Dot1Q PBits          |
| QinQ / TopQ     | TopQ                          | TopQ PBits           |
| QinQ / TopQ     | TopQ BottomQ                  | BottomQ PBits        |
| QinQ / QinQ     | TopQ BottomQ                  | BottomQ PBits        |

**Table 11** Default Dot1P Explicit Marking Actions

| Egress SAP Type | Ingress Packet Preserved Dot1P State                  | Marked (or Remarked) PBits                                   |
|-----------------|-------------------------------------------------------|--------------------------------------------------------------|
| null            | no preserved Dot1P bits                               | none                                                         |
| null            | preserved Dot1P bits                                  | preserved tag PBits remarked using dot1p-value               |
| Dot1Q           | no preserved Dot1P bits                               | new PBits marked using dot1p-value                           |
| Dot1Q           | preserved Dot1P bits                                  | preserved tag PBits remarked using dot1p-value               |
| TopQ            | no preserved Dot1P bits                               | TopQ PBits marked using dot1p-value                          |
| TopQ            | preserved Dot1P bits (used as TopQ and BottomQ PBits) | TopQ PBits marked using dot1p-value, BottomQ PBits preserved |
| QinQ            | no preserved Dot1P bits                               | TopQ PBits and BottomQ PBits marked using dot1p-value        |
| QinQ            | preserved Dot1P bits (used as TopQ and BottomQ PBits) | TopQ PBits and BottomQ PBits marked using dot1p-value        |

**Table 12 QinQ Mark Top Only Explicit Marking Actions**

| Egress SAP Type | Ingress Packet Preserved Dot1P State                  | Marked (or Remarked) PBits                                                      |
|-----------------|-------------------------------------------------------|---------------------------------------------------------------------------------|
| null            | no preserved Dot1P bits                               | none                                                                            |
| null            | preserved Dot1P bits                                  | preserved tag PBits remarked using dot1p-value                                  |
| Dot1Q           | no preserved Dot1P bits                               | new PBits marked using dot1p-value                                              |
| Dot1Q           | preserved Dot1P bits                                  | preserved tag PBits remarked using dot1p-value                                  |
| TopQ            | no preserved Dot1P bits                               | TopQ PBits marked using dot1p-value                                             |
| TopQ            | preserved Dot1P bits (used as TopQ and BottomQ PBits) | TopQ PBits marked using dot1p-value, BottomQ PBits preserved                    |
| QinQ            | no preserved Dot1P bits                               | TopQ PBits marked using dot1p-value, BottomQ PBits marked with zero             |
| QinQ            | preserved Dot1P bits (used as TopQ and BottomQ PBits) | TopQ PBits marked using dot1p-value, BottomQ PBits marked using preserved value |

The QinQ and TopQ SAP PBit/DEI bit marking follows the default behavior defined in [Table 11](#) and [Table 12](#) when **qinq-mark-top-only** is not specified.

The dot1p dot1p-value command must be configured without the qinq-mark-top-only parameter to remove the TopQ PBits only marking restriction.

## agg-rate

**Syntax** [no] **agg-rate**

**Context** config>service>ies>if>sap>egress  
config>service>ies>sub-if>grp-if>sap>egress

**Description** This command is used to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, and **queue-frame-based-accounting**.

---

## rate

|                    |                                                                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>rate</b> <i>rate</i><br><b>no</b> <b>rate</b>                                                                                                                                                                                                  |
| <b>Context</b>     | config>service>ies>if>sap>egress>agg-rate<br>config>service>ies>sub-if>grp-if>sap>egress>agg-rate                                                                                                                                                 |
| <b>Description</b> | This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the agg-rate context to be considered to be active on the context's object (SAP, subscriber, Vport and so on). |
| <b>Parameters</b>  | <i>rate</i> — The <i>rate</i> parameter accepts a value of 1 to 3200000000, or the keyword <b>max</b> .<br><b>Values</b> 1 to 3200000000, <b>max</b>                                                                                              |

## limit-unused-bandwidth

|                    |                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no]</b> <b>limit-unused-bandwidth</b>                                                              |
| <b>Context</b>     | config>service>ies>if>sap>egress>agg-rate<br>config>service>ies>sub-if>grp-if>sap>egress>agg-rate      |
| <b>Description</b> | This command is used to enable (or disable) aggregate rate overrun protection on the agg-rate context. |

## queue-frame-based-accounting

|                    |                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no]</b> <b>queue-frame-based-accounting</b>                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>service>ies>if>sap>egress>agg-rate<br>config>service>ies>sub-if>grp-if>sap>egress>agg-rate                                                                                                                                                                                                                                                                              |
| <b>Description</b> | This command is used to enabled (or disable) frame based accounting on all policers and queues associated with the agg-rate context. Only supported on Ethernet ports. Not supported on HSMDA Ethernet ports. Packet byte offset settings are not included in the applied rate when queue frame-based accounting is configured, but the offsets are applied to the statistics. |

## policer-control-policy

|                    |                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>policer-control-policy</b> <i>policy-name</i><br><b>no</b> <b>policer-control-policy</b> |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap>egress                                                 |
| <b>Description</b> | This command is used to specify a policer control policy to apply to SAP egress.            |



---

|                   |                                                                                             |
|-------------------|---------------------------------------------------------------------------------------------|
| <b>Default</b>    | N/A                                                                                         |
| <b>Parameters</b> | <i>policy-name</i> — Specifies the name of a policer control policy. 32 characters maximum. |

## qinq-mark-top-only

|                    |                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] qinq-mark-top-only</b>                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>service>ies>if>sap>egress                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | When enabled (the encapsulation type of the access port where this SAP is defined as qinq), the <b>qinq-mark-top-only</b> command specifies which P-bits/DEI bit to mark during packet egress. When disabled, both set of P-bits/DEI bit are marked. When the enabled, only the P-bits/DEI bit in the top Q-tag are marked. |
| <b>Default</b>     | no qinq-mark-top-only                                                                                                                                                                                                                                                                                                       |

## qos

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>qos</b> <i>policy-id</i> [ <b>port-redirect-group</b> <i>queue-group-name</i> <b>instance</b> <i>instance-id</i> ]<br><b>no qos</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>service>ies>if>sap>egress<br>config>service>ies>sub-if>grp-if>sap>egress                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command associates a Quality of Service (QoS) policy with an egress Service Access Point (SAP).</p> <p>QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the policy-id does not exist, an error will be returned.</p> <p>The qos command is used to associate both ingress and egress QoS policies. The qos command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.</p> <p>Only one ingress and one egress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a given type will return an error.</p> <p>When an ingress QoS policy is defined on IES ingress IP interface that is bound to a VPLS, the policy becomes associated with every SAP on the VPLS and augments the QoS policy that is defined on each SAP. Packets that are bridged will be processed using the policy defined on the VPLS SAP; packets that are routed will be processed using the policy defined in the IES IP interface-binding context.</p> <p>By default, no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.</p> |

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

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | n/a                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b> | <p><i>policy-id</i> — The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.<br/>1 to 65535</p> <p><i>port-redirect-group</i> — This keyword associates a SAP egress with an instance of a named queue group template on the egress port of a given IOM/IMM/XMA. The <i>queue-group-name</i> and <b>instance</b> <i>instance-id</i> are mandatory parameters when executing the command.</p> <p><i>queue-group-name</i> — Specifies the name of the egress port queue group of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid egress queue group, created under <b>config&gt;port&gt;ethernet&gt;access&gt;egress</b>.</p> <p><b>instance</b> <i>instance-id</i> — Specifies the instance of the named egress port queue group on the IOM/IMM/XMA.</p> <p><b>Values</b> 1 to 40960</p> <p><b>Default</b> 1</p> |

## qos

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>qos</b> <i>policy-id</i> [<b>shared-queuing</b>   <b>multipoint-shared</b>] [<b>fp-redirect-group</b> <i>queue-group-name</i>] [<b>instance</b> <i>instance-id</i>]</p> <p><b>no qos</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | <p>config&gt;service&gt;ies&gt;if&gt;sap&gt;ingress</p> <p>config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;sap&gt;ingress</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP).</p> <p>QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the policy-id does not exist, an error will be returned.</p> <p>The qos command is used to associate both ingress and egress QoS policies. The qos command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.</p> <p>Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type will return an error.</p> <p>By default, no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.</p> |

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

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.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | n/a                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b> | <p><i>policy-id</i> — The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.<br/>1 to 65535</p> <p><b>shared-queuing</b> — Specifies the ingress shared queue policy used by this SAP. When the value of this object is null it means that the SAP will use individual ingress QoS queues instead of the shared ones.</p> <p><b>multipoint-shared</b> — This keyword specifies that this queue-id is for multipoint forwarded traffic only. This queue-id can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. Attempting to map forwarding class unicast traffic to a multipoint queue generates an error; no changes are made to the current unicast traffic queue mapping.</p> <p>A queue must be created as multipoint. The multipoint designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the multipoint keyword, an error is generated and the command will not execute.</p> <p>The multipoint keyword can be entered in the command line on a preexisting multipoint queue to edit queue-id parameters.</p> <p><b>Default</b> Present (the queue is created as non-multipoint).</p> <p><b>Values</b> <b>Multipoint</b> or not present.</p> <p><b>fp-redirect-group</b> — Creates an instance of a named queue group template on the ingress forwarding plane of a given IOM/IMM/XMA. The queue-group-name and instance instance-id are mandatory parameters when executing the command. The named queue group template can contain only policers. If it contains queues, then the command fails.</p> <p><i>queue-group-name</i> — Specifies the name of the queue group template to be instantiated on the forwarding plane of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid ingress queue group template name, configured under <i>config&gt;qos&gt;queue-group-templates</i>.</p> <p><i>instance-id</i> — specifies the instance of the named queue group to be created on the IOM/IMM/XMA ingress forwarding plane.</p> |

## queue-group-redirect-list

|                |                                                                                                   |
|----------------|---------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>queue-group-redirect-list</b> <i>redirect-list-name</i><br><b>no queue-group-redirect-list</b> |
| <b>Context</b> | config>service>ies>if>sap>egress                                                                  |

```
config>service>ies>if>sap>ingress
```

**Description** This command applies a queue group redirect list to the ingress or egress of an interface SAP within an IES or VPRN service. The redirect list is used to redirect traffic to different instances of the default queue group.

This command requires the prior configuration of a default queue group instance, this being the queue group instance specified with the QoS policy under the SAP ingress or egress.

The **no** version of this command removes the queue group redirect list from the SAP.

**Parameters** *redirect-list-name* — Specifies the name of the queue group redirect list up to 32 characters in length.

## queue-override

**Syntax** **[no] queue-override**

**Context** config>service>ies>if>sap>egress  
config>service>ies>if>sap>ingress  
config>service>ies>sub-if>grp-if>sap>egress

**Description** This command enters the context to configure override values for the specified SAP egress QoS queue. These values override the corresponding ones specified in the associated SAP egress or ingress QoS policy.

## hs-secondary-shaper

**Syntax** **hs-secondary-shaper** *policy-name*  
**no hs-secondary-shaper**

**Context** config>service>ies>if>sap>egress>queue-override

**Description** This command configures the HS secondary shaper to be used to apply an aggregate rate and per-scheduling class rates to the SAP egress HSQ queue group.

The **no** form of the command removes the HS secondary shaper override from the configuration returning the SAP egress HSQ queue group to the default HS secondary shaper on that port.

**Parameters** *policy-name* — Specifies the secondary shaper name, up to 32 characters.

## hs-wrr-group

**Syntax** **hs-wrr-group** *group-id* [**create**]  
**hs-wrr-group** *group-id*

---

|                    |                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override                                                                                                            |
| <b>Description</b> | This command configures the egress HS WRR group override parameters.<br><br>The <b>no</b> form of the command removes the group ID from the configuration. |
| <b>Parameters</b>  | <i>group-id</i> — Specifies the HS WRR group ID to override.<br><br><b>Values</b> 1, 2                                                                     |

## class-weight

|                    |                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>class-weight</b> <i>weight</i><br><b>no class-weight</b>                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override>hs-wrr-group                                                                                                                                                                                                                                             |
| <b>Description</b> | This command overrides the class weight of this WRR group at its parent primary shaper relative to the other queues and WRR groups in different HSQ queue groups in the same scheduling class.<br><br>The <b>no</b> form of this command removes the class weight override value from the configuration. |
| <b>Parameters</b>  | <i>weight</i> — Specifies the class weight of the HS WRR group.<br><br><b>Values</b> 1, 2, 4, 8                                                                                                                                                                                                          |

## percent-rate

|                    |                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>percent-rate</b> <i>percent</i><br><b>no percent-rate</b>                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override>hs-wrr-group                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | This command overrides the scheduling rate applied to the HS WRR group as a percentage of the port rate, including both the port's egress rate and port's HS scheduler policy max-rate, if configured. The override rate type must match the corresponding rate type within the applied QoS policy.<br><br>The <b>no</b> form of this command removes the percent rate override value from the configuration. |
| <b>Parameters</b>  | <i>percent</i> — Specifies the percent rate of the HS WRR group.<br><br><b>Values</b> 0.01 to 100.00                                                                                                                                                                                                                                                                                                          |

## rate

|               |                         |
|---------------|-------------------------|
| <b>Syntax</b> | <b>rate</b> <i>rate</i> |
|---------------|-------------------------|

**no rate**

|                    |                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override>hs-wrr-group                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command overrides the scheduling rate applied to the HS WRR group in Kb/s. Alternatively, the keyword <b>max</b> can be specified which removes the bandwidth limitation on the HS WRR group. The override rate type must match the corresponding rate type within the applied QoS policy.</p> <p>The <b>no</b> form of this command removes the rate override value from the configuration.</p> |
| <b>Parameters</b>  | <i>rate</i> — Specifies the scheduling rate of the HS WRR group in Kb/s.                                                                                                                                                                                                                                                                                                                                 |
| <b>Values</b>      | 1 to 2000000000, max                                                                                                                                                                                                                                                                                                                                                                                     |

## queue

|                    |                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>queue</b> <i>queue-id</i> [ <b>create</b> ]<br><b>no queue</b> <i>queue-id</i>                                                                                 |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override<br>config>service>ies>if>sap>ingress>queue-override<br>config>service>ies>sub-if>grp-if>sap>egress>queue-override |
| <b>Description</b> | This command specifies the ID of the queue whose parameters are to be overridden.                                                                                 |
| <b>Parameters</b>  | <i>queue-id</i> — The queue ID whose parameters are to be overridden.                                                                                             |
| <b>Values</b>      | 1 to 32                                                                                                                                                           |

## adaptation-rule

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>adaptation-rule</b> [ <b>pir</b> { <b>max</b>   <b>min</b>   <b>closest</b> }] [ <b>cir</b> { <b>max</b>   <b>min</b>   <b>closest</b> }]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override>queue<br>config>service>ies>if>sap>ingress>queue-override>queue<br>config>service>ies>sub-if>grp-if>sap>egress>queue-override>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command can be used to override specific attributes of the specified queue's adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the queue 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.</p> <p>This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the adaptation rule is performed under the <b>hs-wrr-group</b> within the SAP egress QoS policy.</p> |

The **no** form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no adaptation-rule                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b> | <p><b>pir</b> — The <b>pir</b> parameter defines the constraints enforced when adapting the PIR rate defined within the <b>queue queue-id rate</b> command. The <b>pir</b> parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the <b>rate</b> command is not specified, the default applies.</p> <p><b>max</b> — The <b>max</b> (maximum) option is mutually exclusive with the <b>min</b> and <b>closest</b> options. When <b>max</b> is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the <b>rate</b> command.</p> <p><b>min</b> — The <b>min</b> (minimum) option is mutually exclusive with the <b>max</b> and <b>closest</b> options. When <b>min</b> is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the <b>rate</b> command.</p> <p><b>closest</b> — The <b>closest</b> parameter is mutually exclusive with the <b>min</b> and <b>max</b> parameter. When <b>closest</b> is defined, the operational PIR for the queue will be the rate closest to the rate specified using the <b>rate</b> command.</p> <p><b>cir</b> — The <b>cir</b> parameter defines the constraints enforced when adapting the CIR rate defined within the <b>queue queue-id rate</b> command. The <b>cir</b> parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the <b>cir</b> parameter is not specified, the default constraint applies.</p> |

## avg-frame-overhead

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>avg-frame-overhead percent</b><br><b>no avg-frame-overhead</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override<br>config>service>ies>if>sap>ingress>queue-override>queue<br>config>service>ies>sub-if>grp-if>sap>egress>queue-override>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a Sonet or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for inter-frame gap).</p> <p>When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:</p> |

- **Offered-load** — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- **Frame encapsulation overhead** — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be  $10000 \times 0.1$  or 1000 octets.

For egress Ethernet queues, the frame encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame encapsulation overhead would be  $50 \times 20$  or 1000 octets.

- **Frame based offered-load** — The frame based offered-load is calculated by adding the offered-load to the frame encapsulation overhead. If the offered-load is 10000 octets and the encapsulation overhead was 1000 octets, the frame based offered-load would equal 11000 octets.
- **Packet to frame factor** — The packet to frame factor is calculated by dividing the frame encapsulation overhead by the queue's offered-load (packet based). If the frame encapsulation overhead is 1000 octets and the offered-load is 10000 octets then the packet to frame factor would be  $1000 / 10000$  or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.
- **Frame based CIR** — The frame based CIR is calculated by multiplying the packet to frame factor with the queue's configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame based CIR would be  $500 \times 1.1$  or 550 octets.
- **Frame based within-cir offered-load** — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.

- **Frame based PIR** — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be  $7500 \times 1.1$  or 8250 octets.



- **Frame based within-pir offered-load** — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

**Port scheduler operation using frame transformed rates** — The port scheduler uses the frame based rates to determine the maximum rates that each queue may receive during the within-cir and above-cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

**SAP and subscriber SLA-profile average frame overhead override** — The average frame overhead parameter on a sap-egress may be overridden at an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers an avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

|                   |                                                                                                                                                                                            |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | 0                                                                                                                                                                                          |
| <b>Parameters</b> | <i>percent</i> — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues. |
| <b>Values</b>     | 0 to 100                                                                                                                                                                                   |

## burst-limit

|                    |                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>burst-limit</b> { <b>default</b>   <i>size</i> [ <b>bytes</b>   <b>kilobytes</b> ]}<br><b>no burst-limit</b>                                                                                                               |
| <b>Context</b>     | config>service>ies>sub-if>grp-if>sap>egress>queue-override>queue                                                                                                                                                              |
| <b>Description</b> | The queue <b>burst-limit</b> command defines an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate. |

The **no** form of this command restores the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies. When specified within a queue-override queue context, any current burst limit override for the queue is removed and the queue's burst limit is controlled by its defining policy.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no burst-limit                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b> | <p><b>default</b> — Reverts the queue's burst limit to the system default value.</p> <p><b>size</b> — When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and, by default, is interpreted as the burst limit in kilobytes. If the value is intended to be interpreted in bytes, the <b>bytes</b> qualifier must be added following size.</p> <p><b>Values</b>      0 to 13671 kilobytes<br/>                  0 to or 14000000 bytes</p> <p><b>Default</b>      No default for size; use the <b>default</b> keyword to specify default burst limit.</p> <p><b>bytes</b> — Specifies that the value given for size must be interpreted as the burst limit in bytes.</p> <p><b>kilobytes</b> — Specifies that the value given for size must be interpreted as the burst limit in kilobytes. If neither bytes nor kilobytes is specified, the default qualifier is <b>kilobytes</b>.</p> |

## cbs

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>cbs size-in-kbytes</b><br><b>no cbs</b>                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override>queue<br>config>service>ies>if>sap>ingress>queue-override>queue                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command can be used to override specific attributes of the specified queue's CBS parameters. It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queue's CBS setting into the defined reserved total.</p> |

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.

If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.

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

|                   |                                                                                                                                                                                                                                                                                                 |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no cbs                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b> | <i>size-in-kbytes</i> — The size parameter is an integer number of kilobytes reserved for the queue. For a value of 10 kbytes, enter the number 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimum reserved size can be applied for scheduling purposes). |
| <b>Values</b>     | 0 to 131072, default                                                                                                                                                                                                                                                                            |

## drop-tail

|                    |                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>drop-tail</b>                                                                                                |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override>queue<br>config>service>ies>if>sap>ingress>queue-override>queue |
| <b>Description</b> | This command enters the context to configure queue drop tail parameters.                                        |

## low

|                    |                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>low</b>                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override>queue>drop-tail<br>config>service>ies>if>sap>ingress>queue-override>queue>drop-tail                                                                                   |
| <b>Description</b> | This command enters the context to configure the queue low drop tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets are not accepted into the queue and will be discarded. |

## percent-reduction-from-mbs

|                    |                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>percent-reduction-from-mbs</b> <i>percent</i><br><b>no percent-reduction-from-mbs</b>                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override>queue>drop-tail>low<br>config>service>ies>if>sap>ingress>queue-override>queue>drop-tail>low                                                                                                                                                                                                                                                                 |
| <b>Description</b> | This command overrides the low queue drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and the percentage reduction is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes and out-of-profile packets are not accepted into the queue if its depth is greater than this value, and so will be discarded. |
| <b>Parameters</b>  | <i>percent</i> — Specifies the percentage reduction from the MBS for a queue drop tail.                                                                                                                                                                                                                                                                                                                     |
| <b>Values</b>      | 0 to 100, default                                                                                                                                                                                                                                                                                                                                                                                           |

---

## hs-class-weight

|                    |                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>hs-class-weight</b> <i>weight</i><br><b>no hs-class-weight</b>                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override>queue                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command overrides the class weight of this queue at its parent primary shaper, relative to the other queues and WRR groups in different HSQ queue groups in the same scheduling class.</p> <p>The <b>no</b> form of this command removes the class weight override value from the configuration.</p> |
| <b>Parameters</b>  | <i>weight</i> — Specifies the weight of the queue.<br><b>Values</b> 1, 2, 4, 8                                                                                                                                                                                                                               |

## hs-wred-queue

|                    |                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>hs-wred-queue</b> <i>policy slope-policy-name</i><br><b>no hs-wred-queue</b>                                                                                                                     |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override>queue                                                                                                                                               |
| <b>Description</b> | <p>This command overrides the slope policy applied to the HSQ queue group queue.</p> <p>The <b>no</b> form of this command removes the WRED queue policy override value from the configuration.</p> |
| <b>Parameters</b>  | <i>slope-policy-name</i> — Specifies an existing slope policy name to apply to this HSQ queue group queue.                                                                                          |

## hs-wrr-weight

|                    |                                                                                                                                                                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>hs-wrr-weight</b> <i>weight</i><br><b>no hs-wrr-weight</b>                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override>queue                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command overrides the Weighted Round Robin (WRR) relative weight with which this queue should parent into an HSQ WRR group defined within the associated HS attachment policy.</p> <p>The <b>no</b> form of this command removes the WRR weight override value from the configuration.</p> |
| <b>Parameters</b>  | <i>weight</i> — Specifies the HS WRR group queue weight.<br><b>Values</b> 1 to 127                                                                                                                                                                                                                 |

## mbs

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mbs</b> <i>size</i> {[ <b>bytes</b>   <b>kilobytes</b> ]   <b>default</b> }<br><b>no mbs</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override>queue<br>config>service>ies>if>sap>ingress>queue-override>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command overrides specific attributes of the specified queue's MBS parameters. The MBS is a mechanism to override the default maximum size for the queue.</p> <p>The sum of the MBS for all queues on an egress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.</p> <p>The <b>no</b> form of this command returns the MBS assigned to the queue.</p> |
| <b>Default</b>     | mbs default                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <p><b>size</b> — The size parameter is required when specifying <b>mbs</b> and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional <b>bytes</b> and <b>kilobytes</b> keywords are mutually exclusive and are used to explicitly define whether the size represents bytes or kilobytes.</p> <p><b>Values</b>      0 to 1073741824<br/>                 default</p> <p><b>bytes</b> — When <b>byte</b> is defined, the value given for size is interpreted as the queue's MBS value given in bytes.</p> <p><b>kilobytes</b> — When <b>kilobytes</b> is defined, the value is interpreted as the queue's MBS value given in kilobytes.</p> <p><b>default</b> — Specifying the keyword <b>default</b> sets the MBS to its default value.</p>                                                                                                                                                                                                                                                                     |

## mbs

|                |                                                                                                  |
|----------------|--------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>mbs</b> { <i>size</i> [ <b>kilobytes</b> ]   <b>bytes</b>   <b>default</b> }<br><b>no mbs</b> |
| <b>Context</b> | config>service>ies>if>sap>egress>hsmda-queue-override>queue                                      |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command overrides specific attributes of the specified queue's MBS parameters. The MBS is a mechanism to override the default maximum size for the queue.</p> <p>The sum of the MBS for all queues on an egress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.</p> <p>The <b>no</b> form of this command returns the MBS assigned to the queue.</p> |
| <b>Default</b>     | mbs default                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <p><b>size</b> — Specifies the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kb/s, enter the value 100. A value of 0 causes the queue to discard all packets.</p> <p><b>Values</b>      0 to 2625 kilobytes<br/>                         0 to 2688000 bytes<br/>                         default</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## monitor-depth

|                    |                                                                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] monitor-depth</b>                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | <pre>config&gt;service&gt;ies&gt;if&gt;sap&gt;egress&gt;queue-override&gt;queue config&gt;service&gt;ies&gt;if&gt;sap&gt;ingress&gt;queue-override&gt;queue config&gt;service&gt;vprn&gt;if&gt;sap&gt;egress&gt;queue-override&gt;queue config&gt;service&gt;vprn&gt;if&gt;sap&gt;ingress&gt;queue-override&gt;queue</pre> |
| <b>Description</b> | <p>This command enables queue depth monitoring for the specified queue.</p> <p>The <b>no</b> form of the command removes queue depth monitoring for the specified queue.</p>                                                                                                                                               |

## parent

|                |                                                                                                                                                                                                                                                                                                                                  |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>parent [weight <i>weight</i>] [cir-weight <i>cir-weight</i>]</b>                                                                                                                                                                                                                                                              |
| <b>Context</b> | <pre>config&gt;service&gt;ies&gt;if&gt;sap&gt;egress&gt;queue-override&gt;queue config&gt;service&gt;ies&gt;if&gt;sap&gt;ingress&gt;queue-override&gt;queue config&gt;service&gt;ies&gt;if&gt;sap&gt;egress&gt;sched-override&gt;scheduler config&gt;service&gt;ies&gt;if&gt;sap&gt;ingress&gt;sched-override&gt;scheduler</pre> |

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command can be used to override the scheduler's parent weight and cir-weight information. The weights apply to the associated level/cir-level configured in the applied scheduler policy. The scheduler name must exist in the scheduler policy applied to the ingress or egress of the SAP or multi-service site.</p> <p>The override weights are ignored if the scheduler does not have a parent command configured in the scheduler policy – this allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the SAP/MSS overrides. If the parent scheduler does not exist causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non default weightings for fostered schedulers.</p> <p>The <b>no</b> form of the command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>     | no parent                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><b>weight</b> <i>weight</i> — <b>Weight</b> defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict <b>level</b> defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the queue or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.</p> <p>A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.</p> <p><b>Values</b> 0 to 100</p> <p><b>cir-weight</b> <i>cir-weight</i> — The <b>cir-weight</b> keyword defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same <i>cir-level</i> defined by the <b>cir-level</b> parameter in the applied scheduler policy. Within the strict <b>cir-level</b>, all <b>cir-weight</b> values from active children at that level are summed and the ratio of each active child's <b>cir-weight</b> to the total is used to distribute the available bandwidth at that level. A <b>cir-weight</b> is considered to be active when the policer, queue, or scheduler that the <b>cir-weight</b> pertains to has not reached the CIR and still has packets to transmit.</p> <p>A 0 (zero) <b>cir-weight</b> value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.</p> <p><b>Values</b> 0 to 100</p> |

## percent-rate

|                |                                                                                                    |
|----------------|----------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>percent-rate</b> <i>pir-percent</i> [ <b>cir</b> <i>cir-percent</i> ]<br><b>no percent-rate</b> |
| <b>Context</b> | config>service>ies>if>sap>egress>queue-override>queue                                              |

```
config>service>ies>if>sap>ingress>queue-override>queue
```

**Description** The **percent-rate** command supports a queue's shaping rate and CIR rate as a percentage of the egress port's line rate. When the rates are expressed as a percentage within the template, the actual rate used per instance of the queue group queue-id will vary based on the port speed. For example, when the same template is used to create a queue group on a 1-Gigabit and a 10-Gigabit Ethernet port, the queue's rates will be 10 times greater on the 10 Gigabit port due to the difference in port speeds. This enables the same template to be used on multiple ports without needing to use port based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's shaping and CIR rates will be recalculated based on the defined percentage value.

The rate and percent-rate commands override one another. If the current rate for a queue is defined using the percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at anytime.

An egress port queue group queue rate override may be expressed as either a percentage or an explicit rate independent on how the queue's template rate is expressed.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the **percent-rate** is performed under the **hs-wrr-group** within the SAP egress QoS policy.

The **no** form of this command returns the queue to its default shaping rate and cir rate. When **no percent-rate** is defined within a port egress queue group queue override, the queue reverts to the defined shaping and CIR rates within the egress queue group template associated with the queue.

**Parameters** *pir-percent* — Specifies the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

**Values** 0.01 to 100.00

**Default** 100.00

*cir-percent* — Specifies the queue's committed scheduling rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

**Values** 0.00 to 100.00

**Default** 100.00



## rate

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>rate</b> <i>pir-rate</i> [ <b>cir</b> <i>cir-rate</i> ]<br><b>no rate</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>service>ies>if>sap>egress>queue-override>queue<br>config>service>ies>if>sap>ingress>queue-override>queue<br>config>service>ies>if>sap>egress>sched-override>scheduler<br>config>service>ies>if>sap>ingress>sched-override>scheduler                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.</p> <p>The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.</p> <p>The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile and then out-of-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.</p> <p>The CIR can be used by the queue's parent commands <i>cir-level</i> and <i>cir-weight</i> parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.</p> <p>The <b>rate</b> command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the <i>queue-id</i>.</p> <p>This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the rate is performed under the <b>hs-wrr-group</b> within the SAP egress QoS policy.</p> <p>The <b>no</b> form of the command returns all queues created with the <i>queue-id</i> by association with the QoS policy to the default PIR and CIR parameters (<b>max</b>, 0).</p> |
| <b>Default</b>     | rate max cir 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <p><i>pir-rate</i> — Defines the administrative PIR rate, in kilobits, for the queue. When the <b>rate</b> command is executed, a valid PIR setting must be explicitly defined. When the <b>rate</b> command has not been executed, the default PIR of <b>max</b> is assumed.</p> <p>Fractional values are not allowed and must be given as a positive integer.</p> <p>The actual PIR rate is dependent on the queue's <b>adaptation-rule</b> parameters and the actual hardware where the queue is provisioned.</p> <p>For <b>egress&gt;queue-override&gt;queue</b> and <b>ingress&gt;queue-override&gt;queue</b>:</p> <p><b>Values</b> 1 to 2000000000, <b>max</b> in kb/s</p> <p><b>Default</b> max</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

For **egress>sched-override>scheduler** and **ingress>sched-override>scheduler**:

**Values** 1 to 3200000000, **max** in kb/s

**cir** *cir-rate* — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and must be given as a positive integer. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues.

For **egress>queue-override>queue** and **ingress>queue-override>queue**:

**Values** 0 to 2000000000, **sum**, **max** in kb/s

**Default** 0

For **egress>sched-override>scheduler** and **ingress>sched-override>scheduler**:

**Values** 0 to 3200000000, **sum**, **max** in kb/s

## scheduler-override

|                    |                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] scheduler-override</b>                                                                                                                                                                                                                                |
| <b>Context</b>     | config>service>ies>if>sap>egress<br>config>service>ies>if>sap>ingress                                                                                                                                                                                         |
| <b>Description</b> | This command specifies the set of attributes whose values have been overridden via management on this virtual scheduler. Clearing a given flag will return the corresponding overridden attribute to the value defined on the SAP's ingress scheduler policy. |

## scheduler

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>scheduler scheduler-name [create]</b><br><b>no scheduler scheduler-name</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>if>sap>egress>sched-override<br>config>service>ies>if>sap>ingress>sched-override                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | This command can be used to override specific attributes of the specified scheduler name.<br><br>A scheduler defines a bandwidth controls that limit each child (other schedulers, policers, and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have policers, queues, or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers. |

Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If *scheduler-name* already exists within the policy tier level (regardless of the inclusion of the keyword **create**), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).

If the *scheduler-name* exists within the policy on a different tier (regardless of the inclusion of the keyword **create**), an error occurs and the current CLI context will not change.

If the *scheduler-name* does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:

- Step 1.** The maximum number of schedulers has not been configured.
- Step 2.** The provided *scheduler-name* is valid.
- Step 3.** The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command will not execute, nor will the CLI context change.

If the provided scheduler-name is invalid according to the criteria below, a name syntax error will occur, the command will not execute, and the CLI context will not change.

**Parameters** *scheduler-name* — The name of the scheduler.

**Values** Valid names consist of any string up to 32 characters long 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.

**Default** None. Each scheduler must be explicitly created.

**create** — This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given *scheduler-name*. If the **create** keyword is omitted, **scheduler-name** is not created when the system environment variable **create** is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

## rate

**Syntax** **rate** *pir-rate* [*cir cir-rate*]  
**no rate**

**Context** config>service>ies>if>sap>egress>sched-override>scheduler  
config>service>ies>if>sap>ingress>sched-override>scheduler

```
config>service>ies>if>sap>egress>hsmda-queue-override>queue
config>service>ies>if>sap>ingress>hsmda-queue-override>queue
```

**Description** This command can be used to override specific attributes of the specified scheduler rate. The **rate** command defines the maximum bandwidth that the scheduler can offer its child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child queues and schedulers to operate at their maximum rates.

The **no** form of this command returns the scheduler's PIR and CIR parameters to the value configured in the applied scheduler policy.

**Parameters** *pir-rate* — The **pir** parameter accepts a value of 1 to 3200000000, or the keyword **max**. Any other value will result in an error without modifying the current PIR rate.

**Values** 1 to 3200000000, **max**

**cir** *cir-rate* — This parameter accepts a step-multiplier value that specifies the multiplier used to determine the CIR rate at which the queue will operate. A value of 1 to 3200000000 or the keywords **max** or **sum** is accepted. Any other value will result in an error without modifying the current CIR rate.

If the **cir** is set to **max**, then the CIR rate is set to infinity but is restricted by the PIR rate.

The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers, policers, or queues.

For **egress>sched-override>scheduler** and **ingress>sched-override>scheduler**:

**Values** 0 to 3200000000, **max**, **sum**

## scheduler-policy

**Syntax** **scheduler-policy** *scheduler-policy-name*  
**no scheduler-policy**

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | <pre>config&gt;service&gt;ies&gt;sap&gt;ingress config&gt;service&gt;ies&gt;sap&gt;egress config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;sap&gt;egress config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;sap&gt;ingress config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;sap&gt;egress config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;sap&gt;ingress</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created once the customer site has been appropriately assigned to a chassis port, channel or slot. Scheduler policies are defined in the <b>config&gt;qos&gt;scheduler-policy scheduler-policy-name</b> context.</p> <p>The <b>no</b> form of this command removes the configured ingress or egress scheduler policy from the multi-service customer site. When the policy is removed, the schedulers created due to the policy are removed also making them unavailable for the ingress SAP queues associated with the customer site. Queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler. The SAPs that have ingress queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more queues. When the <b>no scheduler-policy</b> command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.</p> |
| <b>Parameters</b>  | <p><b>scheduler-policy-name:</b> — The <i>scheduler-policy-name</i> parameter applies an existing scheduler policy that was created in the <b>config&gt;qos&gt;scheduler-policy scheduler-policy-name</b> context to create the hierarchy of ingress or egress virtual schedulers. The scheduler names defined within the policy are created and made available to any ingress or egress queues created on associated SAPs.</p> <p><b>Values</b> Any existing valid scheduler policy name.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## 2.5.2.16 ATM Commands

### atm

|                    |                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>atm</b>                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | <pre>config&gt;service&gt;ies&gt;if&gt;sap config&gt;service&gt;ies&gt;sub-if&gt;grp-if&gt;sap</pre>                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command enables access to the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as:</p> <ul style="list-style-type: none"> <li>Configuring ATM port or ATM port-related functionality on MDAs supporting ATM functionality</li> </ul> |

- Configuring ATM-related configuration for ATM-based SAPs that exist on MDAs supporting ATM functionality.

If ATM functionality is not supported for a given context, the command returns an error.

## egress

|                    |                                                                                 |
|--------------------|---------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>egress</b>                                                                   |
| <b>Context</b>     | config>service>ies>if>sap>atm                                                   |
| <b>Description</b> | This command enters the context to configure egress ATM attributes for the SAP. |

## encapsulation

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>encapsulation</b> <i>atm-encap-type</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>service>ies>if>sap>atm                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command configures RFC 2684, <i>Multiprotocol Encapsulation over ATM Adaptation Layer 5</i>, encapsulation for an ATM PVCC delimited SAP.</p> <p>This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition references RFC 2684 and to the ATM Forum LAN Emulation specification. The encapsulation is driven by the services for which the SAP is configured.</p> <p>Ingress traffic that does not match the configured encapsulation will be dropped.</p>                                                                                                                                                                 |
| <b>Default</b>     | encapsulation aal5snap-routed (for IES service SAPs)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <p><i>atm-encap-type</i> — Specifies the encapsulation type.</p> <p><b>Values</b></p> <ul style="list-style-type: none"> <li><b>aal5snap-routed</b> — Routed encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684.</li> <li><b>aal5mux-ip</b> — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684.</li> <li><b>aal5snap-bridged</b> — Bridged encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684.</li> <li><b>aal5mux-bridged-eth-nofcs</b> — Bridged IP encapsulation for VC multiplexed circuit as defined in RFC 2684.</li> </ul> |

## ingress

|               |                |
|---------------|----------------|
| <b>Syntax</b> | <b>ingress</b> |
|---------------|----------------|

---

|                    |                                                             |
|--------------------|-------------------------------------------------------------|
| <b>Context</b>     | config>service>ies>if>sap>atm                               |
| <b>Description</b> | This command configures ingress ATM attributes for the SAP. |

## traffic-desc

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>traffic-desc</b> <i>traffic-desc-profile-id</i><br><b>no traffic-desc</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>service>ies>if>sap>atm>egress<br>config>service>ies>if>sap>atm>ingress                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP).</p> <p>When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction.</p> <p>When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.</p> <p>The <b>no</b> form of the command reverts the traffic descriptor to the default traffic descriptor profile.</p> |
| <b>Default</b>     | The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <i>traffic-desc-profile-id</i> — Specify a defined traffic descriptor profile (see the QoS atm-td-profile command).                                                                                                                                                                                                                                                                                                                                                                                                           |

## oam

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>oam</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>service>ies>if >sap>atm                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command enters the context to configure OAM functionality for a PVCC delimiting a SAP.</p> <p>The ATM-capable MDAs support F5 end-to-end OAM functionality (AIS, RDI, Loopback):</p> <ul style="list-style-type: none"> <li>• ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95</li> <li>• GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996</li> <li>• GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994</li> </ul> |

---

## alarm-cells

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] alarm-cells</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>service>ies>if >sap>atm>oam                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC termination to monitor and report the status of their connection by propagating fault information through the network and by driving PVCC's operational status.</p> <p>When alarm-cells functionality is enabled, a PVCC's operational status is affected when a PVCC goes into an AIS or RDI state because of an AIS/RDI processing (assuming nothing else affects PVCC's operational status, for example, if the PVCC goes DOWN, or enters a fault state and comes back UP, or exits that fault state). RDI cells are generated when PVCC is operationally DOWN. No OAM-specific SNMP trap is raised whenever an endpoint enters/exits an AIS or RDI state, however, if as result of an OAM state change, the PVCC changes operational status, then a trap is expected from an entity the PVCC is associated with (for example a SAP).</p> <p>The <b>no</b> command disables alarm-cells functionality for a PVCC. When alarm-cells functionality is disabled, a PVCC's operational status is no longer affected by a PVCC's OAM state changes due to AIS/RDI processing (when alarm-cells is disabled, a PVCC will change operational status to UP due to alarm-cell processing) and RDI cells are not generated as result of the PVCC going into AIS or RDI state. The PVCC's OAM status, however, will record OAM faults as described above.</p> |
| <b>Default</b>     | Enabled for PVCCs delimiting IES SAPs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## periodic-loopback

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] periodic-loopback</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>service>ies>if >sap>atm>oam                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command enables periodic OAM loopbacks on this SAP. This command is only configurable on IES and VPRN SAPs. When enabled, an ATM OAM loopback cell is transmitted every period as configured in the config&gt;system&gt;atm&gt;oam&gt;loopback-period <i>period</i> context.</p> <p>If a response is not received and consecutive retry-down retries also result in failure, the endpoint will transition to an alarm indication signal/loss of clock state. Then, an ATM OAM loopback cell will be transmitted every period as configured in the loopback-period <i>period</i>. If a response is received for the periodic loopback and consecutive retry-up retries also each receive a response, the endpoint will transition back to the up state.</p> <p>The <b>no</b> form of the command sets the value back to the default.</p> |
| <b>Default</b>     | no periodic-loopback                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |



## 2.5.2.17 IES Interface VRRP Commands

### vrrp

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>        | <b>vrrp</b> <i>virtual-router-id</i> [ <b>owner</b> ] [ <b>passive</b> ]<br><b>no vrrp</b> <i>virtual-router-id</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>       | config>service>ies>if<br>config>service>ies>if>ipv6                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>   | <p>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.</p> <p>Two VRRP nodes can be defined on an IP interface. One, both, or none may be defined as the owner. The nodal context of <b>vrrp</b> <i>virtual-router-id</i> is used to define the configuration parameters for the VRID.</p> <p>The <b>no</b> form of this command removes the specified VRID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the VRID. The VRID does not need to be shutdown in order to remove the virtual router instance.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Default</b>       | n/a                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Special Cases</b> | <p><b>Virtual Router Instance Owner IP Address Conditions</b> — The virtual router instance <b>owner</b> can be created prior to assigning the parent IP interface primary or secondary IP addresses. In this case, the virtual router instance is not associated with an IP address. The operational state of the virtual router instance is down.</p> <p><b>VRRP Owner Command Exclusions</b> — By specifying the VRRP <i>vrid</i> as <b>owner</b>, the following commands are no longer available:</p> <ul style="list-style-type: none"> <li>• <b>vrrp priority</b> — The virtual router instance <b>owner</b> is hard-coded with a <b>priority</b> value of 255 and cannot be changed.</li> <li>• <b>vrrp master-int-inherit</b> — Owner virtual router instances do not accept VRRP advertisement messages; the advertisement interval field is not evaluated and cannot be inherited.</li> <li>• <b>ping-reply</b>, <b>telnet-reply</b> and <b>ssh-reply</b> — The <b>owner</b> virtual router instance always allows Ping, Telnet and SSH if the management and security parameters are configured to accept them on the parent IP interface.</li> <li>• <b>vrrp shutdown</b> — The <b>owner</b> virtual router instance cannot be shut down on the <b>vrrp</b> node. If this was allowed, VRRP messages would not be sent, but the parent IP interface address would continue to respond to ARPs and forward IP packets. Another virtual router instance may detect the missing master due to the termination of VRRP advertisement messages and become master. This would result in two routers responding to ARP requests for the same IP</li> </ul> |

addresses. To shut down the **owner** virtual router instance, use the **shutdown** command in the parent IP interface context. This will prevent VRRP participation, IP ARP reply and IP forwarding. To continue parent IP interface ARP reply and forwarding without VRRP participation, remove the **vrrp vrid** instance.

- **traceroute-reply**

**VRRP Passive Command Exclusions** — By specifying the VRRP *vrid* as **passive**, the following commands related to the master election and processing of VRRP advertisement messages are no longer available:

- **vrrp priority**
- **policy**
- **preempt**
- **master-int-inherit**
- **standby-forwarding**
- **int-delay**
- **message-interval**
- **authentication-key**
- **bfd-enable**

**Parameters** *virtual-router-id* — The virtual-router-id parameter specifies a new virtual router ID or one that can be modified on the IP interface.

**Values** 1 to 255

**owner** — Identifies this virtual router instance as owning the virtual router IP addresses. If the **owner** keyword is not specified at the time of *vrid* creation, the **vrrp backup** commands must be specified to define the virtual router IP addresses. The **owner** keyword is not required when entering the *vrid* for editing purposes. Once created as **owner**, a *vrid* on an IP interface cannot have the **owner** parameter removed. The *vrid* must be deleted, and then recreated without the **owner** keyword, to remove ownership.

**passive** — Identifies this virtual router instance as **passive**, and therefore, owning the virtual router IP addresses. A **passive vrid** does not send or receive VRRP advertisement messages, and is always in either the **master** state (if the interface is operational-up), or the **init** state (if the interface is operational-down). The **passive** keyword is not required when entering the *vrid* for editing purposes. Once a *vrid* on an IP interface is created as **passive**, the parameter cannot be removed from the *vrid*. The *vrid* must be deleted, and then recreated without the **passive** keyword, to remove parameter.

## authentication-key

**Syntax** **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]  
**no authentication-key**

**Context** config>service>ies>if>vrrp

**Description** The **authentication-key** command, within the **vrrp virtual-router-id** context, is used to assign a simple text password authentication key to generate master VRRP advertisement messages and validating 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, the authentication-key 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 is restored to the default value. The authentication-key command may be executed at any time.

To change the current in-use password key on multiple virtual router instances:

- Identify the current master
- Shutdown the virtual router instance on all backups
- Execute the authentication-key command on the master to change the password key
- Execute the authentication-key command and no shutdown command on each backup key

The **no** form of this command restores the default null string to the value of key.

**Default** No default. The authentication data field contains the value 0 in all 16 octets.

**Parameters** *authentication-key* — The *key* parameter identifies the simple text password used when VRRP Authentication Type 1 is enabled on the virtual router instance. Type 1 uses a string eight octets long that is inserted into all transmitted VRRP advertisement messages and compared against all received VRRP advertisement messages. The authentication data fields are used to transmit the key.

The *key* parameter is expressed as a string consisting up to eight alpha-numeric characters. Spaces must be contained in quotation marks (" "). The quotation marks are not considered part of the string.

The string is case sensitive and is left-justified in the VRRP advertisement message authentication data fields. The first field contains the first four characters with the first octet (starting with IETF RFC bit position 0) containing the first character. The second field holds the fifth through eighth characters. Any unspecified portion of the authentication data field is padded with the value 0 in the corresponding octet.

**Values** Any 7-bit printable ASCII character.

|             |                  |          |
|-------------|------------------|----------|
| Exceptions: | Double quote (") | ASCII 34 |
|             | Carriage Return  | ASCII 13 |
|             | Line Feed        | ASCII 10 |
|             | Tab              | ASCII 9  |
|             | Backspace        | ASCII 8  |

*hash-key* — The hash key. The key can be any combination of ASCII characters up to 22 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

**hash** — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

**hash2** — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

## backup

|                    |                                                                        |
|--------------------|------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] backup</b> <i>ip-address</i>                                   |
| <b>Context</b>     | config>service>ies>if>vrrp                                             |
| <b>Description</b> | This command configures virtual router IP addresses for the interface. |

## bfd-enable

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] bfd-enable</b> [ <i>service-id</i> ] <b>interface</b> <i>interface-name</i> <b>dst-ip</b> <i>ip-address</i><br><b>[no] bfd-enable interface</b> <i>interface-name</i> <b>dst-ip</b> <i>ip-address</i> <b>name</b> <i>name</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>service>ies>if>vrrp<br>config>service>ies>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This commands assigns a bi-directional forwarding (BFD) session providing heart-beat mechanism for the given VRRP/SRRP instance. There can be only one BFD session assigned to any given VRRP/SRRP instance, but there can be multiple SRRP/VRRP sessions using the same BFD session.</p> <p>BFD control the state of the associated interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface. The specified interface may not be configured with BFD; however, when it is, the virtual router will then initiate the BFD session.</p> <p>The <b>no</b> form of this command removes BFD from the configuration.</p> |
| <b>Default</b>     | none                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><i>service-id</i> — Specifies the service ID of the interface running BFD.</p> <p><b>Values</b>      <i>service-id</i>: 1 to 2147483648<br/>No service ID indicates a network interface.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

---

**interface** *interface-name* — Specifies the name of the interface running BFD.  
**dst-ip** *ip-address* — Specifies the destination address to be used for the BFD session.  
**name** *name* — Specifies the name, up to 64 characters.

## init-delay

**Syntax** **init-delay** *seconds*  
**no init-delay**

**Context** config>service>ies>if>vrrp

**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

## mac

**Syntax** **mac** *mac-address*  
**no mac**

**Context** config>service>ies>if>vrrp

**Description** This command assigns a specific MAC address to an IES IP interface.  
The **no** form of the command returns the MAC address of the IP interface to the default value.

**Default** The physical MAC address associated with the Ethernet interface that the SAP is configured on (the default MAC address assigned to the interface, assigned by the system).

**Parameters** *mac-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff, where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

## master-int-inherit

**Syntax** [**no**] **master-int-inherit**

**Context** config>service>ies>if>vrrp

**Description** This command allows the master instance to dictate the master down timer (non-owner context only).

---

**Default**    no master-int-inherit

## message-interval

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>message-interval</b> {[ <i>seconds</i> ] [ <b>milliseconds</b> <i>milliseconds</i> ]}<br><b>no message-interval</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>service>ies>if>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>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 than the virtual router instance configured message-interval value will be silently discarded.</p> <p>The message-interval command is available in both non-owner and owner <b>vrrp</b> <i>virtual-router-id</i> nodal contexts. If the message-interval command is not executed, the default message interval of 1 second will be used.</p> <p>The <b>no</b> form of this command restores the default message interval value of 1 second to the virtual router instance.</p> |
| <b>Parameters</b>  | <p><b>seconds</b> — The number of seconds that will transpire before the advertisement timer expires.</p> <p><b>Values</b>        1 to 255</p> <p><b>Default</b>        1</p> <p><b>milliseconds</b> <i>milliseconds</i> — Specifies the time interval, in milliseconds, between sending advertisement messages. This parameter is not supported on non-redundant chassis.</p> <p><b>Values</b>        100 to 900</p>                                                                                                                                                                                                                                                                                                                                                                                 |

## oper-group

|                    |                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>oper-group</b> <i>group-name</i><br><b>no oper-group</b>                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>service>ies>if>vrrp                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command configures VRRP to associate with an operational group. When associated, VRRP notifies the operational group of its state changes so that other protocols can monitor it to provide a redundancy mechanism. When VRRP is the master router (MR), the operational group is up and is down for all other VRRP states.</p> <p>The <b>no</b> form of the command removes the association.</p> |
| <b>Default</b>     | no oper-group                                                                                                                                                                                                                                                                                                                                                                                             |

---

**Parameters**     *group-name* — Specifies the operational group identifier up to 32 characters in length.

## ping-reply

**Syntax**     **ping-reply**  
                 **no ping-reply**

**Context**     config>service>ies>if>vrrp

**Description**     This command enables the non-owner master to reply to ICMP Echo Requests directed at the virtual router instances IP addresses. The ping request can be received on any routed interface.

Ping must not have been disabled at the management security level (either on the parental IP interface or based on the Ping source host address). When ping-reply is not enabled, ICMP Echo Requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to ICMP Echo Requests regardless of the setting of ping-reply configuration.

The ping-reply command is only available in non-owner **vrrp** *virtual-router-id* nodal context. If the ping-reply command is not executed, ICMP Echo Requests to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all ICMP Echo Request messages destined to the non-owner virtual router instance IP addresses.

**Default**     no ping-reply

## policy

**Syntax**     **policy vrrp-policy-id**  
                 **no policy**

**Context**     config>service>ies>if>vrrp

**Description**     This command creates VRRP control policies. The VRRP policy ID must be created by the policy command prior to association with the virtual router instance.

The policy command provides the ability to associate a VRRP priority control policy to a virtual router instance. The policy may be associated with more than one virtual router instance. The priority events within the policy either override or diminish the base-priority dynamically affecting the in-use priority. As priority events clear in the policy, the in-use priority may eventually be restored to the base-priority value.

The policy command is only available in the non-owner **vrrp** *virtual-router-id* nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed by VRRP priority control policies. For non-owner virtual router instances, if the policy command is not executed, the base-priority will be used as the in-use priority.

The **no** form of this command removes any existing VRRP priority control policy association from the virtual router instance. All such associations must be removed prior to the policy being deleted from the system.

**Default** n/a

**Parameters** *vrrp-policy-id* — The vrrp-policy-id parameter associated the corresponding VRRP priority control policy-id with the virtual router instance. The vrrp-policy-id must already exist in the system for the policy command to be successful.

**Values** 1 to 9999

preempt

**Syntax** preempt  
no preempt

**Context** config>service>ies>if>vrrp

**Description** The preempt command provides the ability of overriding an existing non-owner master to the virtual router instance. Enabling preempt mode is almost required for proper operation of the base-priority and vrrp-policy-id definitions on the virtual router instance. If the virtual router cannot preempt an existing non-owner master, the effect of the dynamic changing of the in-use priority is greatly diminished.

The preempt command is only available in the non-owner vrrp virtual-router-id nodal context. The owner may not be preempted due to the fact that the priority of non-owners can never be higher than the owner. The owner will always preempt all other virtual routers when it is available.

Non-owner virtual router instances will only preempt when preempt is set and the current master has an in-use message priority value less than the virtual router instances in-use priority.

A master non-owner virtual router will only allow itself to be preempted when the incoming VRRP Advertisement message Priority field value is one of the following:

- Greater than the virtual router in-use priority value
- Equal to the in-use priority value and the source IP address (primary IP address) is greater than the virtual router instance primary IP address

The **no** form of this command prevents a non-owner virtual router instance from preempting another, less desirable virtual router. Use the preempt command to restore the default mode.

**Default** preempt



## priority

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>priority</b> <i>base-priority</i><br><b>no priority</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>service>ies>if>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>The <b>priority</b> command provides the ability to configure a specific priority value to the virtual router instance. In conjunction with an optional <b>policy</b> command, the <b>base-priority</b> is used to derive the in-use priority of the virtual router instance.</p> <p>The <b>priority</b> command is only available in the non-owner vrrp virtual-router-id nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed. For non-owner virtual router instances, if the <b>priority</b> command is not executed, the <b>base-priority</b> will be set to 100.</p> <p>The <b>no</b> form of this command restores the default value of 100 to <b>base-priority</b>.</p> |
| <b>Parameters</b>  | <p><b>base-priority</b> — The <b>base-priority</b> parameter configures the base priority used by the virtual router instance. If a VRRP Priority Control policy is not also defined, the <b>base-priority</b> will be the in-use priority for the virtual router instance.</p> <p><b>Values</b> 1 to 254</p> <p><b>Default</b> 100</p>                                                                                                                                                                                                                                                                                                                                                                                                    |

## standby-forwarding

|                    |                                                                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] standby-forwarding</b>                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>service>ies>if>vrrp                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command allows the forwarding of packets by a standby router.</p> <p>The <b>no</b> form of the command specifies that a standby router should not forward traffic sent to virtual router's MAC address. However, the standby router should forward traffic sent to the standby router's real MAC address.</p> |
| <b>Default</b>     | no standby-forwarding                                                                                                                                                                                                                                                                                                 |

## ssh-reply

|                |                            |
|----------------|----------------------------|
| <b>Syntax</b>  | <b>[no] ssh-reply</b>      |
| <b>Context</b> | config>service>ies>if>vrrp |

---

**Description** This command enables the non-owner master to reply to SSH Requests directed at the virtual router instances IP addresses. The SSH request can be received on any routed interface. SSH must not have been disabled at the management security level (either on the parental IP interface or based on the SSH source host address). Proper login and CLI command authentication is still enforced.

When ssh-reply is not enabled, SSH packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to SSH regardless of the ssh-reply configuration.

The ssh-reply command is only available in non-owner vrrp virtual-router-id nodal context. If the ssh-reply command is not executed, SSH packets to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all SSH packets destined to the non-owner virtual router instance IP addresses.

**Default** no ssh-reply

## telnet-reply

**Syntax** [no] telnet-reply

**Context** config>service>ies>if>vrrp

**Description** The telnet-reply command enables the non-owner master to reply to TCP port 23 Telnet Requests directed at the virtual router instances IP addresses. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Proper login and CLI command authentication is still enforced.

When telnet-reply is not enabled, TCP port 23 Telnet packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to Telnet Requests regardless of the telnet-reply configuration.

The telnet-reply command is only available in non-owner VRRP nodal context. If the telnet-reply command is not executed, Telnet packets to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all Telnet packets destined to the non-owner virtual router instance IP addresses.

**Default** no telnet-reply

## traceroute-reply

|                    |                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] traceroute-reply</b>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>service>ies>if>vrrp                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.</p> <p>When this command is enabled, a non-owner master can reply to traceroute requests directed to the virtual router instance IP addresses.</p> <p>A non-owner backup virtual router never responds to such traceroute requests regardless of the <b>trace-route-reply</b> status.</p> |
| <b>Default</b>     | no traceroute-reply                                                                                                                                                                                                                                                                                                                                                                                    |

### 2.5.2.18 IPSec Gateway Commands

## ipsec-gw

|                    |                                                   |
|--------------------|---------------------------------------------------|
| <b>Syntax</b>      | <b>ipsec-gw</b> <i>name</i><br><b>no ipsec-gw</b> |
| <b>Context</b>     | config>service>ies>if>sap                         |
| <b>Description</b> | This command configures an IPSec gateway.         |
| <b>Parameters</b>  | <i>name</i> — Specifies IPSec gateway name.       |

## default-secure-service

|                    |                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>default-secure-service</b> <i>service-id</i> <b>interface</b> <i>ip-int-name</i><br><b>default-secure-service</b> <i>name</i> <i>service-name</i> <b>interface</b> <i>ip-int-name</i><br><b>no default-secure-service</b> |
| <b>Context</b>     | config>service>ies>if>sap>ipsec-gateway                                                                                                                                                                                      |
| <b>Description</b> | This command specifies a service ID or service name of the default security service used by this SAP IPSec gateway.                                                                                                          |
| <b>Parameters</b>  | <i>service-id</i> — Specifies a default secure service.                                                                                                                                                                      |

This variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **default-secure-service name** *service-name* variant can be used in all configuration modes.

**Values** {*id* | *svc-name*}

*id*: 1 to 2147483647

*svc-name*: up to 64 characters (*svc-name* is an alias for input only. The *svc-name* gets replaced with an id automatically by SR OS in the configuration).

*service-name* — Identifies the service, up to 64 characters.

## default-tunnel-template

**Syntax** **default-tunnel-template** *ipsec template identifier*  
**no default-tunnel-template**

**Context** config>service>ies>if>sap>ipsec-gateway

**Description** This command configures a default tunnel policy template for the gateway.

## local-gateway-address

**Syntax** **local-gateway-address** *ip-address*  
**no local-gateway-address**

**Context** config>service>ies>if>sap>ipsec-gateway

**Description** This command configures an ipsec-gateway local address.

## pre-shared-key

**Syntax** **pre-shared-key** *key* [*hash* | *hash2*]  
**no pre-shared-key**

**Context** config>service>ies>if>sap>ipsec-gateway

**Description** This command specifies the shared secret between the two peers forming the tunnel.

**Parameters** *key* — Specifies a pre-shared-key for dynamic-keying.

---

## cert

|                    |                                                                     |
|--------------------|---------------------------------------------------------------------|
| <b>Syntax</b>      | <b>cert</b>                                                         |
| <b>Context</b>     | config>service>ies>if>sap>ipsec-gateway                             |
| <b>Description</b> | This command configures cert parameters used by this IPsec gateway. |

## cert-profile

|                    |                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>cert-profile</b> <i>profile-name</i><br><b>no cert-profile</b>                                                                                                       |
| <b>Context</b>     | config>service>ies>if>sap>ipsec-gw>cert                                                                                                                                 |
| <b>Description</b> | This command specifies the cert-profile for the ipsec-tunnel or ipsec-gw. This command will override “cert” and “key” configuration under the ipsec-tunnel or ipsec-gw. |
| <b>Default</b>     | n/a                                                                                                                                                                     |
| <b>Parameters</b>  | <i>profile-name</i> — Specifies the name of cert-profile.                                                                                                               |

## trust-anchor-profile

|                    |                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>trust-anchor</b> <i>trust-anchor-profile-name</i><br><b>no trust-anchor-profile</b>                                                                                                               |
| <b>Context</b>     | config>service>vprn>if>sap>ipsec-gw>cert                                                                                                                                                             |
| <b>Description</b> | This command configures the trust anchor profile name associated with this SAP IPsec tunnel certificate. This command will override “trust-anchor” configuration under the ipsec-tunnel or ipsec-gw. |
| <b>Default</b>     | n/a                                                                                                                                                                                                  |
| <b>Parameters</b>  | <i>trust-anchor-profile-name</i> — Specifies the name of trust-anchor-profile.                                                                                                                       |

## status-verify

|                    |                                                                                          |
|--------------------|------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>status-verify</b>                                                                     |
| <b>Context</b>     | config>service>ies>if>sap>ipsec-gw>cert                                                  |
| <b>Description</b> | This command enters the context to configure certificate status verification parameters. |

---

## default-result

|                    |                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>default-result {revoked   good}</b><br><b>no default-result</b>                                                                                                                                                      |
| <b>Context</b>     | config>service>ies>if>sap>ipsec-gw>cert>status-verify                                                                                                                                                                   |
| <b>Description</b> | This command specifies the default result of Certificate Status Verification (CSV) when both primary and secondary method failed to provide an answer.<br><br>The <b>no</b> form of the command reverts to the default. |
| <b>Default</b>     | default-result revoked                                                                                                                                                                                                  |
| <b>Parameters</b>  | <b>revoked</b> — Specifies that the certificate is considered as revoked.<br><b>good</b> — Specifies that the certificate is considered as acceptable.                                                                  |

## primary

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>primary</b> <i>primary</i> <b>secondary</b> <i>secondary</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>service>ies>if>sap>ipsec-gw>cert>status-verify                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | This command specifies the primary and secondary methods that are used to verify the revocation status of the peer's certificate; either CRL or OCSP. OCSP or CRL uses the corresponding configuration in the CA profile of the issuer of the certificate in question.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>     | primary crl                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <p><i>primary</i> — Specifies the primary method of Certificate Status Verification (CSV) that is used to verify the revocation status of the certificate of the peer.</p> <p><b>Values</b></p> <p>ocsp — Specifies that the OCSP protocol should be used. The OCSP server is configured in the corresponding CA profile.</p> <p>crl — Specifies that the local CRL file should be used. The CRL file is configured in the corresponding CA profile.</p> <p><b>Default</b> crl</p> <p><i>secondary</i> — Specifies the secondary method of CSV that is used to verify revocation status of the certificate of the peer.</p> <p><b>Values</b></p> <p>ocsp — Specifies that the OCSP protocol should be used. The OCSP server is configured in the corresponding CA profile.</p> <p>crl — Specifies that the local CRL file should be used. The CRL file is configured in the corresponding CA profile.</p> |

## local-id

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>local-id type</b> { <b>ipv4</b>   <b>fqdn</b>   <b>ipv6</b> } [ <b>value</b> [ <i>value</i> ]]<br><b>no local-id</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>service>ies>if>sap>ipsec-gateway                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | This command specifies the local ID of the 7750 SR used for IDi or IDr for IKEv2 tunnels. The local-id can only be changed or removed when tunnel or gateway is shutdown.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>     | Depends on local-auth-method such as: <ul style="list-style-type: none"> <li>• Psk:local tunnel ip address</li> <li>• Cert-auth: subject of the local certificate</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <p><b>type</b> — Specifies the type of local ID payload, it could be ipv4 address/FQDN domain name.</p> <p><b>Values</b></p> <ul style="list-style-type: none"> <li>ipv4 — Use ipv4 as the local ID type, the default value is the local tunnel end-point address.</li> <li>ipv6 — Use ipv6 as the local ID type, the default value is the local tunnel end-point address.</li> <li>fqdn — Use FQDN as the local ID type, the value must be configured.</li> <li>dn — Use the subject of the certificate configured for the tunnel or gateway.</li> </ul> <p><i>value</i> — Specifies the value, up to 255 characters.</p> |





## 2.6 IES Show, Clear, and Debug Command Reference

### 2.6.1 Command Hierarchies

- [Show Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)
- [Monitor Commands](#)

#### 2.6.1.1 Show Commands

```
show
 — service
 — customer [customer-id [site customer-site-name [associated-subs]]]
 — egress-label start-label [end-label]
 — id service-id
 — all
 — arp [ip-address] | [mac ieee-address] | [sap sap-id] | [interface ip-int-name] |
 [sdp sdp-id:vc-id] | [summary]
 — arp-host [wholesaler service-id] [sap sap-id | interface interface-name | ip-
 address ip-address[/mask] | mac ieee-address | {[port port-id] [no-inter-
 dest-id | inter-dest-id inter-dest-id]}] [detail]
 — arp-host statistics [sap sap-id | interface interface-name]
 — arp-host summary [interface interface-name | sap]
 — authentication
 — statistics [policy name] [sap sap-id]
 — base [msap]
 — dhcp
 — lease-state [wholesaler service-id] [sap sap-id | sdp sdp-id:vc-id |
 interface interface-name | ip-address ip-address [/mask] | chaddr
 ieee-address | mac iee-address | {[port port-id] [no-inter-dest-id |
 inter-dest-id inter-dest-id]}] [session {none | ipoe}] [detail]
 — statistics [[sap sap-id] | [sdp sdp-id:vc-id] | [interface interface-name]]
 — summary [interface interface-name | saps]
 — gsmp
 — neighbors [group name [ip-address]]
 — sessions [group name]
 — sessions neighbor ip-address port port-number [association |
 statistics]
 — host
 — host-connectivity-verify
 — statistics [summary] [sap sap-id]
 — summary [sap sap-id]
```

- **interface** [{*ip-address* | *ip-int-name*] [*interface-type*] [*detail*] [*family*}] | **summary**
- **retailers**
- **sap** *sap-id* [*detail*]
- **sap** *sap-id* **queue-depth** [*queue queue-id*] [*ingress* | *egress*]
- **sap** *sap-id* **queue-group-redirection** [*ingress* | *egress*]
- **sdp** [*sdp-id* [:*vc-id*]] [*detail*]
- **sdp** **far-end** {*ip-address* | *ipv6-address*} [*detail*]
- **sdp** *sdp-id* [*vc-id*] **l2tpv3**
- **sdp** *sdp-id* [:*vc-id*] **static-isids** [*range-id range-id*]
- **sdp** *sdp-id* [:*vc-id*] **static-isids** **mfib**
- **sdp** *sdp-id* [:*vc-id*] [*detail*] **vccv-bfd** [*session*]
- **sdp** *sdp-id* [:*vc-id*] **mrp**
- **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-prefix/prefix-length*] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**app-profile** *app-profile-name*] [**wholesaler** *service-id*] [**address-origin** *address-origin*] [*detail*]
- **wholesalers**
- **ingress-label** *start-label* [*end-label*]
- **sap-using** [**msap**] [**dyn-script**] [*description*]
- **sap-using** [**sap** *sap-id*] [**vlan-translation** | **anti-spoof**] [*description*]
- **sap-using** {*ingress* | *egress*} **atm-td-profile** *td-profile-id*
- **sap-using** {*ingress* | *egress*} **filter** *any-filter-id*
- **sap-using** {*ingress* | *egress*} **qos-policy** *qos-policy-id* [**msap**]
- **sap-using** **etree**
- **sdp** *sdp-id* **pw-port** [*pw-port-id*]
- **sdp** *sdp-id* **pw-port**
- **sdp** *sdp-id* **pw-port** [*pw-port-id*] [**statistics**]
- **sdp** [**consistent** | **inconsistent** | **na**] **egressifs**
- **sdp** *sdp-id* **keep-alive-history**
- **sdp** **far-end** {*ip-address* | *ipv6-address*} **keep-alive-history**
- **sdp** [*sdp-id*] [*detail*]
- **sdp** **far-end** {*ip-address* | *ipv6-address*} [*detail*]
- **service-using** [**epipe**] [**ies**] [**vppls**] [**vprrn**] [**mirror**] [**apipe**] [**fpipe**] [**ipipe**] [**cpipe**] [**etree**] [**vsd**] [**b-vpls**] [**i-vpls**] [**m-vpls**] [**sdp** *sdp-id*] [**customer** *customer-id*] [**origin** *creation-origin*]
- **subscriber-using** [**service-id** *service-id*] [**sap-id** *sap-id*] [**interface** *ip-int-name*] [**ip** *ip-address[/mask]*] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**app-profile** *app-profile-name*] [**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *intermediate-destination-id*]
- **pw-port** [*pw-port-id*] [*detail*]
- **pw-port** **sdp** [*sdp-id*]
- **pw-port** **sdp** **none**
- **pw-port** *pw-port-id* [**statistics**]
- **router** [*router-instance*]
- **router** **service-name** *service-name*
  - **dhcp**
    - **statistics** [*ip-int-name* | *ip-address*]
    - **summary**
  - **vrrp**
    - **instance**
    - **instance** **interface** *interface-name* [**vrid** *virtual-router-id*]
    - **instance** **interface** *interface-name* **vrid** *virtual-router-id* **ipv6**
    - **statistics**

## 2.6.1.2 Clear Commands

```
clear
 — router
 — dhcp
 — statistics [ip-int-name | ip-address]
 — interface [ip-int-name | ip-address] [icmp] [urpf-stats] [statistics]
 — vrrp
 — interface interface-name [vrid virtual-router-id]
 — interface interface-name vrid virtual-router-id ipv6
 — statistics
 — statistics interface interface-name [vrid virtual-router-id]
 — statistics interface interface-name vrid virtual-router-id ipv6

clear
 — service
 — id service-id
 — arp-host {all | mac ieee-address | sap sap-id | ip-address ip-address[/mask]}
 — arp-host {port port-id | {inter-dest-id intermediate-destination-id | no-inter-dest-id} [port port-id]}
 — arp-host statistics [sap sap-id | interface interface-name]
 — dhcp
 — lease-state ip-address ip-address[/mask] [no-dhcp-release]
 — lease-state mac ieee-address [no-dhcp-release]
 — lease-state sap sap-id [no-dhcp-release]
 — lease-state sdp sdp-id:vc-id [no-dhcp-release]
 — dhcp6
 — lease-state all [no-dhcp-release]
 — lease-state ipv6-address ipv6-prefix [/prefix-length] [no-dhcp-release]
 — lease-state mac ieee-address [no-dhcp-release]
 — lease-state sap sap-id [no-dhcp-release]
 — fdb {all | mac ieee-address | sap sap-id | mesh-sdp sdp-id[:vc-id] | spoke-sdp sdp-id:vc-id}
 — sap sap-id queue-depth [queue queue-id] [ingress | egress]
 — site name
 — spoke-sdp sdp-id:vc-id ingress-vc-label
 — stp
```

## 2.6.1.3 Debug Commands

```
debug
 — router
 — vrrp
 — [no] events
 — events interface ip-int-name [vrid virtual-router-id]
 — events interface ip-int-name vrid virtual-router-id ipv6
 — [no] packets
 — packets interface ip-int-name [vrid virtual-router-id]
 — packets interface ip-int-name vrid virtual-router-id ipv6
 — service
```

- **id** *service-id*
- [no] **arp-host**
  - [no] **ip** *ip-address*
  - [no] **mac** *ieee-address*
  - **mode** {all | dropped-only}
  - [no] **mode**
  - [no] **sap** *sap-id*
- [no] **host-connectivity-verify**
  - [no] **ip** *ip-address*
  - [no] **mac** *ieee-address*
  - [no] **sap** *sap-id*

### 2.6.1.4 Monitor Commands

```
monitor
 — router
 — vrrp
 — instance interface interface-name vr-id virtual-router-id [ipv6] [interval
 seconds] [repeat repeat] [absolute| rate]
```

## 2.6.2 Command Descriptions

- [IES Show Commands](#)
- [IES Clear Commands](#)
- [IES Debug Commands](#)
- [IES Monitor Commands](#)

### 2.6.2.1 IES Show Commands



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

#### customer

**Syntax**     **customer** [*customer-id* [**site** *customer-site-name* [**associated-subs**]]]

**Context**    show>service

**Description** This command displays service customer information.

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

**Values** 1 to 2147483647

**Default** Displays all customer IDs.

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

**associated-subs** — Displays information for the associated subscribers.

**Output** The following output is an example of customer information, and [Table 13](#) 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 13 Show Customer Field Descriptions**

| 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.     |

**Table 13 Show Customer Field Descriptions (Continued)**

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

## egress-label

**Syntax** **egress-label** *start-label* [*end-label*]

**Context** show>service

**Description** This command displays services using the range of ingress labels. If only the mandatory start-label parameter is specified, only services using the specified label are displayed.

If both start-label and end-label parameters are specified, the services using the range of labels X where start-label <= X <= end-label are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

**Parameters** *start-label* — The starting egress label value for which to display services using the label range. If only start-label is specified, services only using start-label are displayed.

**Values** 0, 18432 to 131071

*end-label* — The ending egress label value for which to display services using the label range.

**Values** 18432 to 131071

**Output** The following output is an example of egress label information, and [Table 14](#) 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
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 14** Show Service Egress-Label Field Descriptions

| Label                    | Description                                                                                        |
|--------------------------|----------------------------------------------------------------------------------------------------|
| Svc Id                   | The value that identifies a service.                                                               |
| Sdp Id                   | The value that identifies a SDP.                                                                   |
| Type                     | Indicates whether the SDP binding is a spoke or a mesh.                                            |
| I. Lbl                   | The VC label used by the far-end device to send packets to this device in this service by the SDP. |
| E. Lbl                   | The VC label used by this device to send packets to the far-end device in this service by the SDP. |
| Number of bindings found | The total number of SDP bindings that exist within the specified egress label range.               |



## ingress-label

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ingress-label</b> <i>start-label</i> [ <i>end-label</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | show>service                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command displays services using the range of ingress labels. If only the mandatory <i>start-label</i> parameter is specified, only services using the specified label are displayed.</p> <p>If both <i>start-label</i> and <i>end-label</i> parameters are specified, the services using the range of labels X where <i>start-label</i> &lt;= X &lt;= <i>end-label</i> are displayed.</p> <p>For 7750 only, use the <b>show router vprn-service-id ldp bindings</b> command to display dynamic labels.</p> |
| <b>Parameters</b>  | <p><i>start-label</i> — The starting ingress label value for which to display services using the label range. If only <i>start-label</i> is specified, services only using <i>start-label</i> are displayed.</p> <p><b>Values</b> 0, 2048 to 131071</p> <p><i>end-label</i> — The ending ingress label value for which to display services using the label range.</p> <p><b>Values</b> 2049 to 131071</p> <p><b>Default</b> The <i>start-label</i> value.</p>                                                      |
| <b>Output</b>      | The following output is an example of ingress label information, and <a href="#">Table 15</a> 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
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 15** Show Service Ingress-Label Field Descriptions

| Label                    | Description                                                                                             |
|--------------------------|---------------------------------------------------------------------------------------------------------|
| Svc ID                   | The service identifier.                                                                                 |
| SDP Id                   | The SDP identifier.                                                                                     |
| Type                     | Indicates whether the SDP is spoke or mesh.                                                             |
| I.Lbl                    | The ingress label used by the far-end device to send packets to this device in this service by the SDP. |
| E.Lbl                    | The egress label used by this device to send packets to the far-end device in this service by the SDP.  |
| Number of Bindings Found | The number of SDP bindings within the label range specified.                                            |

## sap-using

**Syntax**

```

sap-using [msap] [dyn-script] [description]
sap-using [sap sap-id] [vlan-translation | anti-spoof] [description]
sap-using [ingress | egress] atm-td-profile td-profile-id
sap-using [ingress | egress] filter any-filter-id
sap-using [ingress | egress] qos-policy qos-policy-id [msap]
sap-using etree

```

**Context** show>service

**Description** Displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs. The optional parameters restrict output to only SAPs matching the specified properties.

**Parameters** **msap** — Displays MSAPs associated to the service. This parameter applies to the 7450 ESS or 7750 SR only.

**dyn-script** — Displays dynamic service SAPs information.

**description** — Displays SAP description.

**sap-id** — Specifies the physical port identifier portion of the SAP definition.

**vlan-translation** — Displays VLAN translation information.

**anti-spoof** — Displays anti-spoof information.

**ingress** — Specifies matching an ingress policy.

**egress** — Specifies matching an egress policy.

**atm-td-profile** *td-profile-id* — Displays SAPs using this traffic description. This parameter only applies to the 7750 SR.

**filter-id** — The ingress or egress filter policy ID for which to display matching SAPs.

**Values** 1 to 65535

**qos-policy-id** — The ingress or egress QoS Policy ID for which to display matching SAPs. This parameter only applies the 7450 ESS and 7750 SR.

**Values** 1 to 65535

**etree** — Specifies matching of SAPs configured as E-Tree SAPs and the corresponding role in the E-Tree services: Leaf-AC, Root-AC or Root-leaf-tag SAPs. SAPs listed as Root-leaf-tag Disabled and Leaf-Ac Disabled function as Root-AC SAPs.

**Output** The following output is an example of SAP information, and [Table 16](#) describes the output fields.

### Sample Output

The following is a sample output for the 7450 ESS:

```
*A:ALA-48# show service sap-using sap 2/1/10:0
=====
Service Access Points Using Port 2/1/10:0
=====
PortId SvcId Ing. Ing. Egr. Egr. Adm Opr
 QoS Fltr QoS Fltr

2/1/10:0 13 1 none 1 none Up Down

Number of SAPs : 1
=====
*A:ALA-48#
```

The following is a sample output for the 7750 SR:

```
*A:ALA-12# show service sap-using egress atm-td-profile 2
=====
Service Access Point Using ATM Traffic Profile 2
=====
PortId SvcId I.QoS I.Fltr E.QoS E.Fltr A.Pol Adm Opr

5/1/1:0/11 511111 2 none 2 none none Up Up
```

```

5/1/1:0/12 511112 2 none 2 none none Up Up
5/1/1:0/13 511113 2 none 2 none none Up Up
5/1/1:0/14 511114 2 none 2 none none Up Up
5/1/1:0/15 511115 2 none 2 none none Up Up
5/1/1:0/16 511116 2 none 2 none none Up Up
5/1/1:0/17 511117 2 none 2 none none Up Up
5/1/1:0/18 511118 2 none 2 none none Up Up
5/1/1:0/19 511119 2 none 2 none none Up Up
5/1/1:0/20 511120 2 none 2 none none Up Up
5/1/1:0/21 511121 2 none 2 none none Up Up
5/1/1:0/22 511122 2 none 2 none none Up Up
5/1/1:0/23 511123 2 none 2 none none Up Up
5/1/1:0/24 511124 2 none 2 none none Up Up
5/1/1:0/25 511125 2 none 2 none none Up Up
...

```

```

=====
*A:ALA-12#

```

**Table 16** Show Service SAP-using Field Descriptions

| 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.                                    |

## sdp-using

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

**Context** show>service

**Description** This command displays services using SDP or far-end address options.

**Parameters** *sdp-id* — Displays only services bound to the specified SDP ID.

**Values** 1 to 17407

*vc-id* — Displays information about the virtual circuit identifier.

**Values** 1 to 4294967295

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

**Default** Services with any far-end IP address.

**Output** The following output is an example of SDP-using information, and [Table 17](#) 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#
```

**Table 17** Show Service Sdp-Using Field Descriptions

| 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.                                                            |
| I.Label    | The label used by the far-end device to send packets to this device in this service by this SDP. |
| E.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**] [**vppls**] [**vprn**] [**mirror**] [**apipe**] [**fpipe**] [**ipipe**] [**cpipe**] [**etree**] [**vsd**] [**b-vppls**] [**i-vppls**] [**m-vppls**] [**sdp sdp-id**] [**customer customer-id**] [**origin creation-origin**]

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | show>service                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | This command displays the services matching certain usage properties. If no optional parameters are specified, all services defined on the system are displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <p><b>epipe</b> — Displays matching Epipe services (applies only to the 7450 ESS and 7750 SR).</p> <p><b>ies</b> — Displays matching IES instances.</p> <p><b>vpls</b> — Displays matching VPLS instances (applies only to the 7450 ESS and 7750 SR).</p> <p><b>vprn</b> — Displays matching VPRN services.</p> <p><b>mirror</b> — Displays mirror services.</p> <p><b>apipe</b> — Displays matching Apipe services. (applies only to the 7450 ESS and 7750 SR).</p> <p><b>fpipe</b> — Displays matching Fpipe services (applies only to the 7450 ESS and 7750 SR).</p> <p><b>ipipe</b> — Displays matching Ipipe services (applies only to the 7450 ESS and 7750 SR).</p> <p><b>cpipe</b> — Displays matching Cpipe services.</p> <p><b>etree</b> — Displays etree service information.</p> <p><b>vsd</b> — Displays VSD service information.</p> <p><b>b-vpls</b> — Displays matching B-VPLS services.</p> <p><b>i-vpls</b> — Displays matching I-VPLS services.</p> <p><b>m-vpls</b> — Displays matching M-VPLS services.</p> <p><b>sdp-id</b> — Displays only services bound to the specified SDP ID.</p> <p><b>Values</b> 1 to 17407</p> <p><b>Default</b> Services bound to any SDP ID.</p> <p><b>customer-id</b> — Displays services only associated with the specified customer ID.</p> <p><b>Values</b> 1 to 2147483647</p> <p><b>Default</b> Services associated with a customer.</p> <p><b>creation-origin</b> — Specifies the method by which the service was created.</p> <p><b>Values</b> manual, multi-segment-p-w, dyn-script, vsd</p> |
| <b>Output</b>      | The following output is an example of service-using information, and <a href="#">Table 18</a> 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 18** Show Service Service-Using Field Descriptions

| 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. |

## subscriber-using

**Syntax** **subscriber-using** [**service-id** *service-id*] [**sap-id** *sap-id*] [**interface** *ip-int-name*] [**ip** *ip-address[/mask]*] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**app-profile** *app-profile-name*] [**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *intermediate-destination-id*]

**Context** show>service

**Description** This command displays subscribers using certain options.

**Parameters** *service-id* — Display subscriber information about the specified service ID.

**Values** 1 to 2147483648, *svc-name*: 64 characters max

*sap-id* — Specifies the physical port identifier portion of the SAP definition.

*ip-int-name* — Displays subscriber information about the specified interface. 32 characters maximum.

*ip-address[/mask]* — Displays subscriber information about the specified IP address.

**Values** ip-address: a.b.c.d  
mask: 1 to 32

*ieee-address* — Displays subscriber information about the specified MAC address.

**Values**      *xx:xx:xx:xx:xx:xx* or *xx-xx-xx-xx-xx-xx* (cannot be all zeros)

*sub-profile-name* — Displays subscriber information about the specified subscriber profile name. 32 characters maximum.

*sla-profile-name* — Displays subscriber information about the specified SLA profile name. 32 characters maximum.

*app-profile-name* — Specifies the application profile name. 32 characters maximum.

*port-id* — Specifies the port ID.

**no-inter-dest-id** — Displays the information about no intermediate destination ID.

*intermediate-destination-id* — Displays information about the specified intermediate destination ID. 32 characters maximum.

id

|                    |                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>id</b> <i>service-id</i>                                                                                                                                         |
| <b>Context</b>     | show>service                                                                                                                                                        |
| <b>Description</b> | This command displays information for a particular service-id.                                                                                                      |
| <b>Parameters</b>  | <i>service-id</i> — The unique service identification number to identify the service in the service domain. 64 characters maximum.<br><b>Values</b> 1 to 2148007980 |

all

|                    |                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>all</b>                                                                                                              |
| <b>Context</b>     | show>service>id                                                                                                         |
| <b>Description</b> | This command displays detailed information for all aspects of the service.                                              |
| <b>Output</b>      | The following output is an example of service id information, and <a href="#">Table 19</a> describes the output fields. |

**Sample Output**

```
A:PE# show service id 500 all

=====
Service Detailed Information
=====
Service Id : 500 Vpn Id : 0
Service Type : IES
Name : 500
```



```

Description : (Not Specified)
Customer Id : 1 Creation Origin : manual
Last Status Change: 08/31/2018 16:07:55
Last Mgmt Change : 08/31/2018 16:07:55
Admin State : Up Oper State : Up
SAP Count : 2 SDP Bind Count : 0

ETH-CFM service specifics

Tunnel Faults : ignore

Service Destination Points(SDPs)

No Matching Entries

Service Access Points

SAP 1/1/9:1

Service Id : 500
SAP : 1/1/9:1 Encap : q-tag
Description : (Not Specified)
Admin State : Up Oper State : Up
Flags : None
Multi Svc Site : None
Last Status Change: 08/31/2018 16:07:55
Last Mgmt Change : 08/31/2018 16:07:55
Sub Type : regular
Dot1Q Ethertype : 0x8100 QinQ Ethertype : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU : 1518 Oper MTU : 1518
Ingr IP Fltr-Id : n/a Egr IP Fltr-Id : n/a
Ingr Mac Fltr-Id : n/a Egr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : both
Egr Agg Rate Limit : max
Q Frame-Based Acct : Disabled Limit Unused BW : Disabled
Qinq-vlan-translation : None Qinq-vlan-translation Ids : None

Acct. Pol : None Collect Stats : Disabled

Anti Spoofing : None Dynamic Hosts : Enabled
Avl Static Hosts : 0 Tot Static Hosts : 0
Calling-Station-Id : n/a

Application Profile: None
Transit Policy : None
AARP Id : None

Oper Group : (none) Monitor Oper Grp : (none)
Host Lockout Plcy : n/a
Lag Link Map Prof : (none)

```

Bandwidth : Not-Applicable  
Oper DCpu Prot Pol\*: \_default-access-policy

-----  
ETH-CFM SAP specifics  
-----

Tunnel Faults : n/a AIS : Disabled  
MC Prop-Hold-Timer : n/a  
Squelch Levels : None  
Collect Lmm Stats : Disabled  
LMM FC Stats : None  
LMM FC In Prof : None

-----  
QOS  
-----

Ingress qos-policy : 1 Egress qos-policy : 1  
Ingress FP QGrp : (none) Egress Port QGrp : (none)  
Ing FP QGrp Inst : (none) Egr Port QGrp Inst: (none)  
Shared Q plcy : n/a Multipoint shared : Disabled  
I. Sched Pol : (Not Specified)  
E. Sched Pol : (Not Specified)  
I. Policer Ctl Pol : (Not Specified)  
E. Policer Ctl Pol : (Not Specified)  
E. HS Sec. Shaper : (Not Specified)  
I. QGrp Redir. List: (Not Specified)  
E. QGrp Redir. List: (Not Specified)

-----  
Sap Aggregate Stats  
-----

|                         | Packets | Octets |
|-------------------------|---------|--------|
| Ingress                 |         |        |
| Aggregate Offered : 0   |         | 0      |
| Aggregate Forwarded : 0 |         | 0      |
| Aggregate Dropped : 0   |         | 0      |
| Egress                  |         |        |
| Aggregate Forwarded : 0 |         | 0      |
| Aggregate Dropped : 0   |         | 0      |

-----  
Sap Statistics  
-----

Last Cleared Time : 08/31/2018 16:19:29

|                                      | Packets | Octets |
|--------------------------------------|---------|--------|
| CPM Ingress : 0                      |         | 0      |
| Forwarding Engine Stats              |         |        |
| Dropped : 0                          |         | 0      |
| Received Valid : 0                   |         | 0      |
| Off. HiPrio : 0                      |         | 0      |
| Off. LowPrio : 0                     |         | 0      |
| Off. Uncolor : 0                     |         | 0      |
| Off. Managed : 0                     |         | 0      |
| Queueing Stats(Ingress QoS Policy 1) |         |        |
| Dro. HiPrio : 0                      |         | 0      |
| Dro. LowPrio : 0                     |         | 0      |

```

For. InProf : 0 0
For. OutProf : 0 0

Queueing Stats(Egress QoS Policy 1)
Dro. In/InplusProf : 0 0
Dro. Out/ExcProf : 0 0
For. In/InplusProf : 0 0
For. Out/ExcProf : 0 0

Sap per Queue stats

 Packets Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio : 0 0
Off. LowPrio : 0 0
Dro. HiPrio : 0 0
Dro. LowPrio : 0 0
For. InProf : 0 0
For. OutProf : 0 0

Egress Queue 1
For. In/InplusProf : 0 0
For. Out/ExcProf : 0 0
Dro. In/InplusProf : 0 0
Dro. Out/ExcProf : 0 0
* indicates that the corresponding row element may have been truncated.

SAP pw-1:500

Service Id : 500
SAP : pw-1:500 Encap : q-tag
Description : (Not Specified)
Admin State : Up Oper State : Up
Flags : None
Multi Svc Site : None
Last Status Change : 08/31/2018 16:09:16
Last Mgmt Change : 08/31/2018 16:07:55
Sub Type : regular
Split Horizon Group : (Not Specified)

Admin MTU : 1518 Oper MTU : 8660
Ingr IP Fltr-Id : n/a Egr IP Fltr-Id : n/a
Ingr Mac Fltr-Id : n/a Egr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : both
Egr Agg Rate Limit : max

Limit Unused BW : Disabled
Qinq-vlan-translation : None Qinq-vlan-translation Ids : None

Acct. Pol : None Collect Stats : Disabled

Anti Spoofing : None Dynamic Hosts : Enabled
Avl Static Hosts : 0 Tot Static Hosts : 0
Calling-Station-Id : n/a

Application Profile: None

```

```

Transit Policy : None
AARP Id : None

Oper Group : (none) Monitor Oper Grp : (none)
Host Lockout Plcy : n/a
Lag Link Map Prof : (none)
Bandwidth : Not-Applicable
Oper DCpu Prot Pol*: _default-access-policy

```

-----  
ETH-CFM SAP specifics  
-----

```

Tunnel Faults : n/a AIS : Disabled
MC Prop-Hold-Timer : n/a
Squelch Levels : None
Collect Lmm Stats : Disabled
LMM FC Stats : None
LMM FC In Prof : None

```

-----  
QOS  
-----

```

Ingress qos-policy : 1 Egress qos-policy : 500
Ingress FP QGrp : (none) Egress Port QGrp : (none)
Ing FP QGrp Inst : (none) Egr Port QGrp Inst: (none)
Shared Q plcy : n/a Multipoint shared : Disabled
I. Sched Pol : (Not Specified)
E. Sched Pol : pw-sap
I. Policer Ctl Pol : (Not Specified)
E. Policer Ctl Pol : (Not Specified)
E. HS Sec. Shaper : (Not Specified)
I. QGrp Redir. List: (Not Specified)
E. QGrp Redir. List: (Not Specified)

```

-----  
Sap Aggregate Stats  
-----

|                     | Packets | Octets |
|---------------------|---------|--------|
| Ingress             |         |        |
| Aggregate Offered   | : 0     | 0      |
| Aggregate Forwarded | : 0     | 0      |
| Aggregate Dropped   | : 0     | 0      |
| Egress              |         |        |
| Aggregate Forwarded | : 0     | 0      |
| Aggregate Dropped   | : 0     | 0      |

-----  
Sap Statistics  
-----

Last Cleared Time : 08/31/2018 16:19:29

|             | Packets | Octets |
|-------------|---------|--------|
| CPM Ingress | : 0     | 0      |

Forwarding Engine Stats

|                |     |   |
|----------------|-----|---|
| Dropped        | : 0 | 0 |
| Received Valid | : 0 | 0 |
| Off. HiPrio    | : 0 | 0 |
| Off. LowPrio   | : 0 | 0 |

```
Off. Uncolor : 0 0
Off. Managed : 0 0
```

Queueing Stats(Ingress QoS Policy 1)

```
Dro. HiPrio : 0 0
Dro. LowPrio : 0 0
For. InProf : 0 0
For. OutProf : 0 0
```

Queueing Stats(Egress QoS Policy 500)

```
Dro. In/InplusProf : 0 0
Dro. Out/ExcProf : 0 0
For. In/InplusProf : 0 0
For. Out/ExcProf : 0 0
```

Sap per Queue stats

```

Packets Octets
```

Ingress Queue 1 (Unicast) (Priority)

```
Off. HiPrio : 0 0
Off. LowPrio : 0 0
Dro. HiPrio : 0 0
Dro. LowPrio : 0 0
For. InProf : 0 0
For. OutProf : 0 0
```

Egress Queue 1

```
For. In/InplusProf : 0 0
For. Out/ExcProf : 0 0
Dro. In/InplusProf : 0 0
Dro. Out/ExcProf : 0 0
```

\* indicates that the corresponding row element may have been truncated.

Service Interfaces

Interface

```

If Name : int
Admin State : Up Oper (v4/v6) : Up/Down
Down Reason V6 : ifProtoOperDown
Protocols : None
IP Addr/mask : 192.168.0.1/24 Address Type : Primary
IGP Inhibit : Disabled Broadcast Address : Host-ones
HoldUp-Time : 0 Track Srrp Inst : 0
Description : N/A
Ignore Port State : None
```

Details

```

Description : (Not Specified)
If Index : 4 Virt. If Index : 4
Last Oper Chg : 08/31/2018 16:14:20 Global If Index : 258
Mon Oper Grp : None
Srrp En Rtng : Disabled Hold time : N/A
SAP Id : pw-1:500
TOS Marking : Untrusted If Type : IES
```

---

|                                                   |                     |                    |            |
|---------------------------------------------------|---------------------|--------------------|------------|
| SNTP B.Cast                                       | : False             | IES ID             | : 500      |
| MAC Address                                       | : d8:2c:ff:00:00:00 | Mac Accounting     | : Disabled |
| Ingress stats                                     | : Disabled          | IPv6 DAD           | : Enabled  |
| TCP MSS V4                                        | : 0                 | TCP MSS V6         | : 0        |
| 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 |
| ARP Learn Unsolic*                                | : Disabled          | ND Learn Unsolici* | : None     |
| ARP Proactive Ref*                                | : Disabled          | ND Proactive Refr* | : None     |
| ARP Populate Host                                 | : Disabled          | ND Populate Host   | : Disabled |
| ARP Route Tag                                     | : 0                 | ND Route tag       | : 0        |
| IP MTU                                            | : (default)         |                    |            |
| IP Oper MTU                                       | : 8660              | Host Conn Verify   | : Disabled |
| ARP Populate                                      | : Disabled          | Cflowd (multicast) | : None     |
| SHCV pol IPv4                                     | : None              |                    |            |
| Cflowd (unicast)                                  | : None              |                    |            |
| 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          |                    |            |
| uRPF Select VPRN                                  | : False             |                    |            |
| PTP HW Assist                                     | : Disabled          |                    |            |
| Rx Pkts                                           | : 0                 | Rx Bytes           | : 0        |
| Rx V4 Pkts                                        | : N/A               | Rx V4 Bytes        | : N/A      |
| Rx V6 Pkts                                        | : N/A               | Rx V6 Bytes        | : N/A      |
| Tx DBcast Dis. Pk*                                | : 0                 | Tx DBcast Dis. By* | : 0        |
| Mpls Rx Pkts                                      | : 0                 | Mpls Rx Bytes      | : 0        |
| Mpls Tx Pkts                                      | : 0                 | Mpls Tx Bytes      | : 0        |
| OperDCpuProtPlcy                                  | : N/A               |                    |            |
| Proxy ARP Details                                 |                     |                    |            |
| Rem Proxy ARP                                     | : Disabled          | Local Proxy ARP    | : Disabled |
| Policies                                          | : none              |                    |            |
| Proxy Neighbor Discovery Details                  |                     |                    |            |
| Local Pxy ND                                      | : Disabled          |                    |            |
| Policies                                          | : none              |                    |            |
| DHCP no local server                              |                     |                    |            |
| DHCP Details                                      |                     |                    |            |
| Description                                       | : (Not Specified)   |                    |            |
| Admin State                                       | : Down              | Lease Populate     | : 0        |
| Gi-Addr                                           | : 192.168.0.1*      | Gi-Addr as Src Ip  | : Disabled |
| * = inferred gi-address from interface IP address |                     |                    |            |
| Action                                            | : Keep              | Trusted            | : Disabled |
| DHCP Proxy Details                                |                     |                    |            |
| Admin State                                       | : Down              |                    |            |
| Lease Time                                        | : N/A               |                    |            |
| Emul. Server                                      | : Not configured    |                    |            |

DHCP Relay Proxy Details

Relay ucast : none  
Siaddr ovr : N/A

Subscriber Authentication Details

Auth Policy : None

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

ISA Tunnel redundant next-hop information

Static Next-Hop :  
Dynamic Next-Hop :

ICMP Details

|                   |                |                |      |
|-------------------|----------------|----------------|------|
| Redirects         | : Number - 100 | Time (seconds) | - 10 |
| Unreachables      | : Number - 100 | Time (seconds) | - 10 |
| TTL Expired       | : Number - 100 | Time (seconds) | - 10 |
| Parameter Problem | : Number - 100 | Time (seconds) | - 10 |
| ICMP Mask Reply   | : True         |                |      |

ICMPv6 Details

|                   |                |                |      |
|-------------------|----------------|----------------|------|
| Packet Too Big    | : Number - 100 | Time (seconds) | - 10 |
| Parameter Problem | : Number - 100 | Time (seconds) | - 10 |
| Redirects         | : Number - 100 | Time (seconds) | - 10 |
| Time Exceeded     | : Number - 100 | Time (seconds) | - 10 |
| Unreachables      | : Number - 100 | Time (seconds) | - 10 |

IPCP Address Extension Details

Peer IP Addr : Not configured  
Peer Pri DNS Addr : Not configured  
Peer Sec DNS Addr : Not configured

Network Domains Associated

default

-----  
Admin Groups

-----  
No Matching Entries  
-----

-----  
Srlg Groups

-----  
No Matching Entries  
-----

-----  
QoS Queue-Group Redirection Details

-----  
Ingress FP QGrp : (none) Egress Port QGrp : (none)

Ing FP QGrp Inst : (none)

Egr Port QGrp Inst: (none)

-----  
Interface  
-----

|                   |                   |                   |             |
|-------------------|-------------------|-------------------|-------------|
| If Name           | : int1            |                   |             |
| Admin State       | : Up              | Oper (v4/v6)      | : Up/Down   |
| Down Reason V6    | : ifProtoOperDown |                   |             |
| Protocols         | : None            |                   |             |
| IP Addr/mask      | : 20.0.0.1/30     | Address Type      | : Primary   |
| IGP Inhibit       | : Disabled        | Broadcast Address | : Host-ones |
| HoldUp-Time       | : 0               | Track Srrp Inst   | : 0         |
| Description       | : N/A             |                   |             |
| Ignore Port State | : None            |                   |             |

-----  
Details  
-----

|                     |                       |                    |            |
|---------------------|-----------------------|--------------------|------------|
| Description         | : (Not Specified)     |                    |            |
| If Index            | : 5                   | Virt. If Index     | : 5        |
| Last Oper Chg       | : 08/31/2018 16:07:55 | Global If Index    | : 259      |
| Mon Oper Grp        | : None                |                    |            |
| Srrp En Rtng        | : Disabled            | Hold time          | : N/A      |
| SAP Id              | : 1/1/9:1             |                    |            |
| TOS Marking         | : Untrusted           | If Type            | : IES      |
| SNTP B.Cast         | : False               | IES ID             | : 500      |
| MAC Address         | : d8:2c:01:01:00:09   | Mac Accounting     | : Disabled |
| Ingress stats       | : Disabled            | IPv6 DAD           | : Enabled  |
| TCP MSS V4          | : 0                   | TCP MSS V6         | : 0        |
| 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 |
| ARP Learn Unsolic*  | : Disabled            | ND Learn Unsolici* | : None     |
| ARP Proactive Refr* | : Disabled            | ND Proactive Refr* | : None     |
| ARP Populate Host   | : Disabled            | ND Populate Host   | : Disabled |
| ARP Route Tag       | : 0                   | ND Route tag       | : 0        |
| IP MTU              | : (default)           |                    |            |
| IP Oper MTU         | : 1500                |                    |            |
| ARP Populate        | : Disabled            | Host Conn Verify   | : Disabled |
| SHCV pol IPv4       | : None                |                    |            |
| Cflowd (unicast)    | : None                | Cflowd (multicast) | : None     |
| 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            |                    |            |
| uRPF Select VPRN    | : False               |                    |            |
| PTP HW Assist       | : Disabled            |                    |            |
| Rx Pkts             | : 3                   | Rx Bytes           | : 0        |
| Rx V4 Pkts          | : N/A                 | Rx V4 Bytes        | : N/A      |
| Rx V6 Pkts          | : N/A                 | Rx V6 Bytes        | : N/A      |
| Tx Pkts             | : 5                   | Tx Bytes           | : 530      |
| Tx V4 Pkts          | : 5                   | Tx V4 Bytes        | : 530      |
| Tx V4 Discard Pkts  | : 0                   | Tx V4 Discard Byt* | : 0        |
| Tx V6 Pkts          | : 0                   | Tx V6 Bytes        | : 0        |



---

```

Tx V6 Discard Pkts: 0
Tx DBcast Dis. Pk*: 0
Mpls Rx Pkts : 0
Mpls Tx Pkts : 0
OperDCpuProtPlcy : N/A

Tx V6 Discard Byt*: 0
Tx DBcast Dis. By*: 0
Mpls Rx Bytes : 0
Mpls Tx Bytes : 0

Proxy ARP Details
Rem Proxy ARP : Disabled
Policies : none

Local Proxy ARP : Disabled

Proxy Neighbor Discovery Details
Local Pxy ND : Disabled
Policies : none

DHCP no local server

DHCP Details
Description : (Not Specified)
Admin State : Down
Gi-Addr : 20.0.0.1*
* = inferred gi-address from interface IP address
Lease Populate : 0
Gi-Addr as Src Ip : Disabled

Action : Keep
Trusted : Disabled

DHCP Proxy Details
Admin State : Down
Lease Time : N/A
Emul. Server : Not configured

DHCP Relay Proxy Details
Relay ucast : none
Siaddr ovr : N/A

Subscriber Authentication Details
Auth Policy : None

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

ISA Tunnel redundant next-hop information
Static Next-Hop :
Dynamic Next-Hop :

ICMP Details
Redirects : Number - 100
Unreachables : Number - 100
TTL Expired : Number - 100
Parameter Problem: Number - 100
ICMP Mask Reply : True
Time (seconds) - 10
Time (seconds) - 10
Time (seconds) - 10
Time (seconds) - 10

ICMPv6 Details

```

```

Packet Too Big : Number - 100 Time (seconds) - 10
Parameter Problem: Number - 100 Time (seconds) - 10
Redirects : Number - 100 Time (seconds) - 10
Time Exceeded : Number - 100 Time (seconds) - 10
Unreachables : Number - 100 Time (seconds) - 10

```

```

IPCP Address Extension Details
Peer IP Addr : Not configured
Peer Pri DNS Addr : Not configured
Peer Sec DNS Addr : Not configured

```

```

Network Domains Associated
default

```

```

Admin Groups

```

```

No Matching Entries

```

```

Srlg Groups

```

```

No Matching Entries

```

```

QoS Queue-Group Redirection Details

```

```

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

```

```

=====
* indicates that the corresponding row element may have been truncated.

```

```

A:PE#

```

**Table 19**      **Show All Service-ID Field Descriptions**

| 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.                                          |

**Table 19 Show All Service-ID Field Descriptions (Continued)**

| Label                             | Description (Continued)                                                                                                                                          |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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: GRE or MPLS.                                                                                                     |
| Admin State                       | The administrative state of this SDP.                                                                                                                            |
| Oper State                        | The operational state of this SDP.                                                                                                                               |
| Ingress Label                     | The label used by the far-end device to send packets to this device in this service by this SDP.                                                                 |
| Egress Label                      | The label used by this device to send packets to the far-end device in this service by this SDP.                                                                 |
| Ingress Filter                    | The ID of the ingress filter policy.                                                                                                                             |
| Egress Filter                     | The ID of the egress filter policy.                                                                                                                              |
| Far End                           | Specifies the IP address of the remote end of the GRE or 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.                                                 |

**Table 19 Show All Service-ID Field Descriptions (Continued)**

| Label                  | Description (Continued)                                                                                                                                                                                                                                                              |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. If the SDP type is GRE, then the following message displays:<br>“SDP Delivery Mechanism is not MPLS”. |
| Number of SDPs         | The total number SDPs applied to this service ID.                                                                                                                                                                                                                                    |
| Service Access Points  |                                                                                                                                                                                                                                                                                      |
| Service Id             | The service identifier.                                                                                                                                                                                                                                                              |
| Port Id                | The ID of the access port where this SAP is defined.                                                                                                                                                                                                                                 |
| Description            | Generic information about the SAP.                                                                                                                                                                                                                                                   |
| Encap                  | The value of the label used to identify this SAP on the access port.                                                                                                                                                                                                                 |
| Admin State            | The desired state of the SAP.                                                                                                                                                                                                                                                        |
| Oper State             | The operating state of the SAP.                                                                                                                                                                                                                                                      |
| Last Changed           | The date and time of the last change.                                                                                                                                                                                                                                                |
| Admin MTU              | The largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.                                                                                                                            |
| Oper MTU               | The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.                                                                                                                     |
| Ingress qos-policy     | The SAP ingress QoS policy ID.                                                                                                                                                                                                                                                       |
| Egress qos-policy      | The SAP egress QoS policy ID.                                                                                                                                                                                                                                                        |
| Ingress Filter-Id      | The SAP ingress filter policy ID.                                                                                                                                                                                                                                                    |
| Egress Filter-Id       | The SAP egress filter policy ID.                                                                                                                                                                                                                                                     |
| Multi Svc Site         | Indicates the multi-service site that the SAP is a member.                                                                                                                                                                                                                           |
| Ingress sched-policy   | Indicates the ingress QoS scheduler for the SAP.                                                                                                                                                                                                                                     |
| 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.                                                                                                                                                                                                    |

**Table 19 Show All Service-ID Field Descriptions (Continued)**

| Label                          | Description (Continued)                                                                                                                                     |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Offered Low Priority           | The number of low priority packets, as determined by the SAP ingress QoS policy.                                                                            |
| Forwarded In Profile           | The number of in-profile packets or octets (rate below CIR) forwarded.                                                                                      |
| Forwarded Out Profile          | The number of out-of-profile packets or octets (rate above CIR) forwarded.                                                                                  |
| Queueing Stats                 |                                                                                                                                                             |
| 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.                                                                                  |
| SAP per Queue stats            |                                                                                                                                                             |
| Ingress Queue 1                | The index of the ingress QoS queue of this SAP.                                                                                                             |
| High priority offered          | The packets or octets count of the high priority traffic for the SAP.                                                                                       |
| High priority dropped          | The number of high priority traffic packets/octets dropped.                                                                                                 |
| Low priority offered           | The packets or octets count of the low priority traffic.                                                                                                    |
| Low priority dropped           | The number of low priority traffic packets/octets dropped.                                                                                                  |
| In profile forwarded           | The number of in-profile packets or octets (rate below CIR) forwarded.                                                                                      |
| Out profile forwarded          | The number of out-of-profile octets (rate above CIR) forwarded.                                                                                             |
| Egress Queue 1                 | The index of the egress QoS queue of the SAP.                                                                                                               |
| In profile forwarded           | The number of in-profile packets or octets (rate below CIR) forwarded.                                                                                      |
| IPCP Address Extension Details |                                                                                                                                                             |
| In profile dropped             | The number of in-profile packets or octets dropped for the SAP.                                                                                             |
| Peer IP Addr                   | Specifies the remote IP address to be assigned to the far-end of the associated PPP/MLPPP link via IPCP extensions.                                         |
| Peer Pri DNS Addr              | Specifies a unicast IPv4 address for the primary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions.              |
| Peer Sec DNS Addr              | Specifies a unicast IPv4 address for the secondary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions (optional). |

## arp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>arp</b> [ <i>ip-address</i> ]   [ <b>mac</b> <i>ieee-address</i> ]   [ <b>sap</b> <i>sap-id</i> ]   [ <b>interface</b> <i>ip-int-name</i> ]   [ <b>sdp</b> <i>sdp-id:vc-id</i> ]   [ <b>summary</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | show>service>id                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | Displays the ARP table for the IES instance. The ARP entries are displayed uniquely. Each MAC associated with the child group-interfaces are displayed with each ARP entry. They do not reflect actual ARP entries but are displayed along the interfaces ARP entry for easy lookup.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><i>ieee-address</i> — Displays only ARP entries in the ARP table with the specified 48-bit MAC address. The MAC address can be expressed in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> where <i>aa</i>, <i>bb</i>, <i>cc</i>, <i>dd</i>, <i>ee</i> and <i>ff</i> are hexadecimal numbers.</p> <p><b>Values</b>      <i>xx:xx:xx:xx:xx:xx</i> or <i>xx-xx-xx-xx-xx-xx</i></p> <p><i>sap-id</i> — Displays SAP information for the specified SAP ID.</p> <p><b>interface</b> — Specifies matching service ARP entries associated with the IP interface.</p> <p><i>ip-address</i> — The IP address of the interface for which to display matching ARP entries.</p> <p><b>Values</b>      <i>a.b.c.d</i></p> <p><i>ip-int-name</i> — The IP interface name for which to display matching ARPs. 32 characters maximum.</p> <p><i>sdp-id</i> — The SDP identifier.</p> <p><b>Values</b>      1 to 17407</p> <p><i>vc-id</i> — The virtual circuit identifier.</p> <p><b>Values</b>      1 to 4294967295</p> |
| <b>Output</b>      | <p><b>Show Service-ID ARP</b></p> <p>The following output is an example of service ID ARP information, and <a href="#">Table 20</a> describes the output fields.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## Sample Output

```
A:ALA-49# show service id 88 arp
=====
ARP Table
=====
```

| IP Address | MAC Address       | Type    | Expiry    | Interface | SAP       |
|------------|-------------------|---------|-----------|-----------|-----------|
| 10.30.1.1  | 76:1e:ff:00:01:41 | Other   | 00h00m00s | ies30     | lag-1:30  |
| 10.31.1.1  | 76:1e:ff:00:01:41 | Other   | 00h00m00s | ies30     | lag-1:30  |
| 10.37.1.1  | 00:00:00:00:00:00 | Other   | 00h00m00s | foo2      | n/a       |
| 10.20.1.1  | 76:1e:ff:00:00:00 | Other   | 00h00m00s | s2        | subscrib* |
|            | 76:1e:ff:00:01:41 |         |           | g3        | lag-1     |
| 10.20.1.10 | 00:00:aa:aa:aa:dd | Managed | 00h00m00s | g3        | lag-1:11  |
| 10.20.1.11 | 00:00:aa:aa:aa:dd | Managed | 00h00m00s | g3        | lag-1:11  |

```

10.20.1.12 00:00:aa:aa:aa:dd Managed 00h00m00s g3 lag-1:11
10.38.1.1 76:1e:ff:00:00:00 Other 00h00m00s s3 subscrib*
 76:21:04:01:00:01 g5 4/1/1
 76:21:04:01:00:01 g7 4/1/1
10.39.1.1 76:1e:ff:00:00:00 Other 00h00m00s s3 subscrib*
 76:21:04:01:00:01 g5 41/1/1
 76:21:04:01:00:01 g7 4/1/1
10.38.1.2 76:22:07:01:00:01 Managed 00h00m00s g7 4/1/1:25*
10.38.10.1 76:22:07:01:00:01 Managed 00h00m00s g7 4/1/1:25*
10.38.99.1 76:22:07:01:00:01 Managed 00h00m00s g7 4/1/1:25*

```

=====

\* indicates that the corresponding row element may have been truncated.

A:ALA-49#

**Table 20**      **Show Service-ID ARP Field Descriptions**

| Label       | Description                                                                                                                                                                                                    |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP Address  | The IP address.                                                                                                                                                                                                |
| MAC Address | The specified MAC address.                                                                                                                                                                                     |
| Type        | Static — FDB entries created by management.<br>Learned — Dynamic entries created by the learning process.<br>OAM — Entries created by the OAM process.<br>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.                                                                                                                                                                                                    |

## arp-host

**Syntax**      **arp-host** [*wholesaler service-id*] [**sap** *sap-id* | **interface** *interface-name* | **ip-address** *ip-address*[/*mask*] | **mac** *ieee-address* | {[**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *inter-dest-id*]}] [**detail**]

**arp-host statistics** [**sap** *sap-id* | **interface** *interface-name*]

**arp-host summary** [**interface** *interface-name* | **saps**]

**Context**      show>service>id

**Description**      This command displays ARP host related information.

**Parameters**      *service-id* — Specifies the wholesaler service ID.

**Values**          1 to 2148278316

*sap-id* — Displays SAP information for the specified SAP ID.

*interface-name* — Displays information for the specified IP interface. 32 characters maximum.

*ip-address[/mask]* — Shows information for the specified IP address and mask.

**Values**      a.b.c.d.  
                 mask: 1 to 32

*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.

*port-id* — Specifies the port ID.

**no-inter-dest-id** — Displays the information about no intermediate destination ID.

*inter-dest-id* — Displays information about the specified intermediate destination ID.

**saps** —

**Output**      The following output is an example of ARP host information.

### Sample Output

```
*A:Dut-C# show service id 2 arp-host
=====
ARP host table, service 2
=====
IP Address Mac Address Sap Id Remaining MC
 Time Stdbby

192.168.1.2 00:80:00:00:00:01 2/1/5:2 00h04m41s
192.168.1.3 00:80:00:00:00:02 2/1/5:2 00h04m42s
192.168.1.4 00:80:00:00:00:03 2/1/5:2 00h04m43s
192.168.1.5 00:80:00:00:00:04 2/1/5:2 00h04m44s
192.168.1.6 00:80:00:00:00:05 2/1/5:2 00h04m45s
192.168.1.7 00:80:00:00:00:06 2/1/5:2 00h04m46s
192.168.1.8 00:80:00:00:00:07 2/1/5:2 00h04m47s
192.168.1.9 00:80:00:00:00:08 2/1/5:2 00h04m48s
192.168.1.10 00:80:00:00:00:09 2/1/5:2 00h04m49s
192.168.1.11 00:80:00:00:00:0a 2/1/5:2 00h04m50s

Number of ARP hosts : 10
=====
*A:Dut-C#

*A:Dut-C# show service id 2 arp-host ip-address 192.168.1.2 detail
=====
ARP hosts for service 2
=====
Service ID : 2
IP Address : 192.168.1.2
MAC Address : 00:80:00:00:00:01
SAP : 2/1/5:2
Remaining Time : 00h04m58s
```



```

Sub-Ident : "alu_1_2"
Sub-Profile-String : ""
SLA-Profile-String : ""
App-Profile-String : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
RADIUS-User-Name : "192.168.1.2"

Session Timeout (s) : 301
Start Time : 02/09/2009 16:35:07
Last Auth : 02/09/2009 16:36:34
Last Refresh : 02/09/2009 16:36:38
Persistence Key : N/A

Number of ARP hosts : 1
=====
*A:Dut-C#

*A:Dut-C# show service id 2 arp-host statistics
=====
ARP host statistics
=====
Num Active Hosts : 20
Received Triggers : 70
Ignored Triggers : 10
Ignored Triggers (overload) : 0
SHCV Checks Forced : 0
Hosts Created : 20
Hosts Updated : 40
Hosts Deleted : 0
Authentication Requests Sent : 40
=====
*A:Dut-C#

*A:Dut-C# show service id 2 arp-host summary
=====
ARP host Summary, service 2
=====
Sap Used Provided Admin State

sap:2/1/5:2 20 8000 inService

Number of SAPs : 1

=====
*A:Dut-C#

```

## statistics

|                    |                                                              |
|--------------------|--------------------------------------------------------------|
| <b>Syntax</b>      | <b>statistics [policy name] [sap sap-id]</b>                 |
| <b>Context</b>     | show>service>id>authentication                               |
| <b>Description</b> | Displays session authentication statistics for this service. |

- Parameters**    *name* — Specifies the subscriber authentication policy statistics to display.  
                   *sap-id* — Specifies the SAP ID statistics to display.
- Output**        The following output is an example of authentication statistics information.

**Sample Output**

```
*A:ALA-1# show service id 11 authentication statistics

Authentication statistics

Interface / SAP Authentication Authentication
 Successful Failed

abc-11-10.1.0.254 1582 3

Number of entries: 1
=====
*A:ALA-1#
```

## authentication

- Syntax**        **authentication**
- Context**       **show>service>id**
- Description**   This command enables the context to display subscriber authentication information.

## base

- Syntax**        **base [msap]**
- Context**       **show>service>id**
- Description**   This command displays basic information about this IES service.
- Parameters**    **msap** — Displays MSAPs associated to the service.
- Output**        The following output is an example of basic IES service information.

**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#

```

## router

- Syntax** **router** [*router-instance*]  
**router service-name service-name**
- Context** show
- Description** This command displays router instance information.
- Parameters** *router-instance* — Specifies the router name or service ID used to specify the router instance.
- Values**
- router-instance router-name | vprn-svc-id*
- |             |                                                   |
|-------------|---------------------------------------------------|
| router-name | Base   management   cpm-vr-name   vpls-management |
|             | Default: Base                                     |
| vprn-svc-id | 1 to 2147483647                                   |
| cpm-vr-name | 32 characters maximum                             |
- service-name* — Specifies the service name used to identify the router instance. 64 characters maximum.

## dhcp

- Syntax** **dhcp**
- Context** show>service>id
- Description** This command enables the context to display DHCP information for the specified service.

## lease-state

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lease-state</b> [ <b>wholesaler</b> <i>service-id</i> ] [ <b>sap</b> <i>sap-id</i>   <b>sdp</b> <i>sdp-id:vc-id</i>   <b>interface</b> <i>interface-name</i>   <b>ip-address</b> <i>ip-address</i> [/mask]   <b>chaddr</b> <i>ieee-address</i>   <b>mac</b> <i>ieee-address</i>   {[ <b>port</b> <i>port-id</i> ] [ <b>no-inter-dest-id</b>   <b>inter-dest-id</b> <i>inter-dest-id</i> ]}] [ <b>session</b> { <b>none</b>   <b>ipoe</b> }] [ <b>detail</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | show>service>id>dhcp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | This command displays DHCP lease state related information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <p><i>service-id</i> — Specifies the service ID of the wholesaler.</p> <p><b>Values</b> 1 to 2148278316, <i>svc-name: 64 characters maximum</i></p> <p><i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition.</p> <p><i>sdp-id</i> — The SDP identifier.</p> <p><b>Values</b> 1 to 17407</p> <p><i>vc-id</i> — The virtual circuit ID on the SDP ID for which to display information.</p> <p><b>Values</b> 1 to 4294967295</p> <p><i>interface-name</i> — Displays information for the specified IP interface.</p> <p><i>ip-address</i>[/mask] — Displays information associated with the specified IP address.</p> <p><i>ieee-address</i> — Specifies the source MAC address.</p> <p><b>Values</b> xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx (cannot be all zeros)</p> <p><i>port-id</i> — Specifies the port ID.</p> <p><b>no-inter-dest-id</b> — Displays the information about no intermediate destination ID.</p> <p><i>inter-dest-id</i> — Displays information about the specified intermediate destination ID. 32 characters maximum.</p> <p><b>session</b> — Shows DHCPv4 lease states for hosts that are associated with an IPoE session or for hosts that are not associated with an IPoE session.</p> <p><b>Values</b> none, ipoe</p> <p><b>detail</b> — Displays detailed information.</p> |
| <b>Output</b>      | The following is a sample output of the <b>dhcp</b> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

### Sample Output

```
A:ALA-_Dut-A# show service id 13 dhcp lease-state
=====
DHCP lease state table, service 13
=====
IP Address Mac Address Sap/Sdp Id Remaining Lease MC
 LifeTime Origin Stdby

10.13.40.1 00:00:00:00:00:13 1/1/1:13 00h00m58s Radius

```

```

Number of lease states : 1
=====
A:ALA-_Dut-A#

A:ALA-_Dut-A# show service id 13 dhcp lease-state detail
=====
DHCP lease states for service 13
=====
Service ID : 13
IP Address : 10.13.40.1
Mac Address : 00:00:00:00:00:13
Interface : ies-13-10.13.1.1
SAP : 1/1/1:13
Remaining Lifetime : 00h00m58s
Persistence Key : N/A

Sub-Ident : "TEST"
Sub-Profile-String : "ADSL GO"
SLA-Profile-String : "BE-Video"
Lease ANCP-String : ""

Sub-Ident origin : Radius
Strings origin : Radius
Lease Info origin : Radius

Ip-Netmask : 255.255.0.0
Broadcast-Ip-Addr : 10.13.255.255
Default-Router : N/A
Primary-Dns : 10.13.254.254
Secondary-Dns : 10.13.254.253

ServerLeaseStart : 12/24/2006 23:44:07
ServerLastRenew : 12/24/2006 23:44:07
ServerLeaseEnd : 12/24/2006 23:45:07
Session-Timeout : 0d 00:01:00
DHCP Server Addr : N/A

Persistent Relay Agent Information
 Circuit Id : ancstb6_Dut-A|13|ies-13-10.13.1.1|0|13
 Remote Id : stringtest

Number of lease states : 1
=====
A:ALA-_Dut-A#

```

### Routed CO Output Example

```

A:ALA-_Dut-A# show service id 13 dhcp lease-state
=====
DHCP lease state table, service 13
=====
IP Address Mac Address Sap/Sdp Id Remaining Lease MC
 LifeTime Origin Stdby

10.13.40.1 00:00:00:00:00:13 1/1/1:13 00h00m58s Radius

Number of lease states : 1

```

```

=====
A:ALA-_Dut-A#

A:ALA-_Dut-A# show service id 13 dhcp lease-state detail
=====
DHCP lease states for service 13
=====
Service ID : 13
IP Address : 10.13.40.1
Mac Address : 00:00:00:00:00:13
Subscriber-interface : ies-13-10.13.1.1
Group-interface : intf-13
SAP : 1/1/1:13
Remaining Lifetime : 00h00m58s
Persistence Key : N/A

Sub-Ident : "TEST"
Sub-Profile-String : "ADSL GO"
SLA-Profile-String : "BE-Video"
Lease ANCP-String : ""

Sub-Ident origin : Radius
Strings origin : Radius
Lease Info origin : Radius

Ip-Netmask : 255.255.0.0
Broadcast-Ip-Addr : 10.13.255.255
Default-Router : N/A
Primary-Dns : 10.13.254.254
Secondary-Dns : 10.13.254.253

ServerLeaseStart : 12/24/2006 23:48:23
ServerLastRenew : 12/24/2006 23:48:23
ServerLeaseEnd : 12/24/2006 23:49:23
Session-Timeout : 0d 00:01:00
DHCP Server Addr : N/A

Persistent Relay Agent Information
 Circuit Id : ancstb6_Dut-A|13|intf-13|0|13
 Remote Id : stringtest

Number of lease states : 1
=====
A:ALA-_Dut-A#

```

### Wholesaler/Retailer Output Example

```

A:ALA-_Dut-A# show service id 2000 dhcp lease-state detail
=====
DHCP lease states for service 2000

Wholesaler 1000 Leases

Service ID : 1000
IP Address : 10.13.1.254

```

```

Mac Address : 00:00:00:00:00:13
Subscriber-interface : whole-sub
Group-interface : intf-13
Retailer : 2000
Retailer If : retail-sub
SAP : 1/1/1:13
Remaining Lifetime : 00h09m59s
Persistence Key : N/A

Sub-Ident : "TEST"
Sub-Profile-String : "ADSL GO"
SLA-Profile-String : "BE-Video"
Lease ANCP-String : ""

Sub-Ident origin : Retail DHCP
Strings origin : Retail DHCP
Lease Info origin : Retail DHCP

Ip-Netmask : 255.255.0.0
Broadcast-Ip-Addr : 10.13.255.255
Default-Router : N/A
Primary-Dns : N/A
Secondary-Dns : N/A

ServerLeaseStart : 12/25/2006 00:29:41
ServerLastRenew : 12/25/2006 00:29:41
ServerLeaseEnd : 12/25/2006 00:39:41
Session-Timeout : 0d 00:10:00
DHCP Server Addr : 10.232.237.2

Persistent Relay Agent Information
 Circuit Id : 1/1/1:13
 Remote Id : stringtest

Number of lease states : 1
=====
A:ALA- _Dut-A#

```

## statistics

**Syntax** **statistics** **[[sap sap-id] | [sdp sdp-id:vc-id] | [interface interface-name]**

**Context** show>service>id>dhcp

**Description** Displays DHCP statistics information.

**Parameters**

*sap-id* — Specifies the physical port identifier portion of the SAP definition.

*sdp-id* — The SDP identifier.

**Values** 1 to 17407

*vc-id* — The virtual circuit ID on the SDP ID for which to display information.

**Values** 1 to 4294967295

*interface-name* — Displays information for the specified IP interface. 32 characters maximum.

## summary

|                    |                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>summary</b> [ <b>interface</b> <i>interface-name</i>   <b>saps</b> ]                                                               |
| <b>Context</b>     | show>service>id>dhcp                                                                                                                  |
| <b>Description</b> | Displays DHCP configuration summary information.                                                                                      |
| <b>Parameters</b>  | <p><i>interface-name</i> — Specifies the interface name. 32 characters maximum.</p> <p><b>saps</b> — Displays SAPs per interface.</p> |
| <b>Output</b>      | The following output is an example of DHCP summary information, and <a href="#">Table 21</a> describes the output fields.             |

### Sample Output (7750 SR)

```
A:ALA-49# show service id 88 dhcp summary
=====
DHCP Summary, service 88
=====
Interface Name Arp Used/ Info Admin
 SapId/Sdp Populate Provided Option State

Sector A No 0/0 Keep Up
 sap:7/1/1.2.2 0/0
grp-if No 0/1 Keep Down
 sap:2/2/2:0 0/1
 sap:2/2/2:0 0/1
test No 0/0 Keep Up
 sap:10/1/2:0 0/0

Interfaces: 3
=====
A:ALA-49#
```

### Sample Output (7450 ESS)

```
A:ALA-49# show service id 88 dhcp summary
=====
DHCP Summary, service 88
=====
Interface Name Arp Used/ Info Admin
 SapId/Sdp Populate Provided Option State

Sector A No 0/0 Keep Up
 sap:7/1/1.2.2 0/0
 sap:2/2/2:0 0/1
test No 0/0 Keep Up
 sap:10/1/2:0 0/0

```



```
Interfaces: 3
=====
A:ALA-49#
```

**Table 21 Show DHCP Summary Field Descriptions**

| Label          | Description                                                                                                                                                                                                                                                                                                           |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface Name | Name of the router interface.                                                                                                                                                                                                                                                                                         |
| Arp Populate   | Specifies whether or not ARP populate is enabled.                                                                                                                                                                                                                                                                     |
| Used/Provided  | Used — The number of lease-states that are currently in use on a specific interface, that is, the number of clients on that interface got an IP address by DHCP. This value is always less than or equal to the 'Provided' field.<br>Provided — The lease-populate value that is configured for a specific interface. |
| Info Option    | Indicates whether Option 82 processing is enabled on the interface.                                                                                                                                                                                                                                                   |
| Admin State    | Indicates the administrative state.                                                                                                                                                                                                                                                                                   |

## gsmp

**Syntax** **gsmp**

**Context** show>service>id

**Description** This command enables the context to display GSMP information.

## neighbors

**Syntax** **neighbors** [**group** *name* [*ip-address*]]

**Context** show>service>id>gsmp

**Description** This command displays GSMP neighbor information.

**Parameters** **group** — A GSMP group defines a set of GSMP neighbors which have the same properties.  
*name* — Specifies a GSMP group name is unique only within the scope of the service in which it is defined. 32 characters maximum.  
*ip-address* — Specifies the ip-address of the neighbor.

**Values** a.b.c.d (unicast address only)

**Output** These commands show the configured neighbors per service, regardless of the fact there exists an open TCP connection with this neighbor. The admin state is shown because for a neighbor to be admin enabled, the service, gsmp node, group node and the neighbor node in this service must all be in 'no shutdown' state. Session gives the number of session (open TCP connections) for each configured neighbor.

### Sample Output

```
A:active>show>service>id>gsmp# neighbors
=====
GSMP neighbors
=====
Group Neighbor AdminState Sessions

dslaml 192.168.1.2 Enabled 0
dslaml 192.168.1.3 Enabled 0

Number of neighbors shown: 2
=====
A:active>show>service>id>gsmp#

A:active>show>service>id>gsmp# neighbors group dslaml
=====
GSMP neighbors
=====
Group Neighbor AdminState Sessions

dslaml 192.168.1.2 Enabled 0
dslaml 192.168.1.3 Enabled 0

Number of neighbors shown: 2
=====
A:active>show>service>id>gsmp#
A:active>show>service>id>gsmp# neighbors group dslaml 192.168.1.2
=====
GSMP neighbors
=====
Group Neighbor AdminState Sessions

dslaml 192.168.1.2 Enabled 0
=====
A:active>show>service>id>gsmp#
```

## sessions

**Syntax** **sessions** [*group name*]  
**sessions neighbor** *ip-address port port-number* [**association** | **statistics**]

**Context** show>service>id>gsmp

**Description** This command displays GSMP sessions information.

- Parameters**
- group** — A GSMP group defines a set of GSMP neighbors which have the same properties.
  - name** — Specifies a GSMP group name is unique only within the scope of the service in which it is defined. 32 characters maximum.
  - ip-address** — Specifies the ip-address of the neighbor.
    - Values** a.b.c.d (unicast address only)
  - port-number** — Specifies the neighbor TCP port number use for this ANCP session.
    - Values** 0 to 65535
  - association** — Displays to what object the ANCP-string is associated.
  - statistics** — Displays statistics information about an ANCP session known to the system.
- Output** The following output gives information about the open TCP connections with DSLAMs.

### Sample Output

```
A:active>show>service>id>gsmp# sessions
```

```
=====
GSMP sessions for service 999 (VPRN)
=====
Port Ngbr-IPAddr Gsmp-Group

40590 192.168.1.2 dslam1

Number of GSMP sessions : 1
=====
A:active>show>service>id>gsmp#
```

```
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590
```

```
=====
GSMP sessions for service 999 (VPRN), neighbor 192.168.1.2, Port 40590
=====
State : Established
Peer Instance : 1 Sender Instance : a3cf58
Peer Port : 0 Sender Port : 0
Peer Name : 12:12:12:12:12:12 Sender Name : 00:00:00:00:00:00
timeouts : 0 Max. Timeouts : 3
Peer Timer : 100 Sender Timer : 100
Capabilities : DTD OAM
Conf Capabilities : DTD OAM
Priority Marking : dscp nc2
Local Addr. : 192.168.1.4
Conf Local Addr. : N/A
=====
A:active>show>service>id>gsmp#
```

```
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 association
```

```
=====
ANCP-Strings
=====
```

```
ANCP-String Assoc. State

No ANCP-Strings found
=====
A:active>show>service>id>gsmp#

A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 statistics
=====
GSMP session stats, service 999 (VPRN), neighbor 192.168.1.2, Port 40590
=====
Event Received Transmitted

Dropped 0 0
Syn 1 1
Syn Ack 1 1
Ack 14 14
Rst Ack 0 0
Port Up 0 0
Port Down 0 0
OAM Loopback 0 0
=====
A:active>show>service>id>gsmp#
```

The association command gives an overview of each ANCP string received from this session (applies only to the 7450 ESS and 7750 SR).

```
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 association
=====
ANCP-Strings
=====
ANCP-String Assoc.
State

7330-ISAM-E47 atm 1/1/01/01:19425.64048 ANCP Up

Number of ANCP-Strings : 1
=====
A:active>show>service>id>gsmp#
```

host

|             |                                                                                                                   |
|-------------|-------------------------------------------------------------------------------------------------------------------|
| Syntax      | host                                                                                                              |
| Context     | show>service>id                                                                                                   |
| Description | Displays static hosts configured for this IES service.                                                            |
| Output      | The following output is an example of host information, and <a href="#">Table 22</a> describes the output fields. |

Sample Output

```
*A:ALA-48# show service id 88 host
```

```
=====
Static Hosts for service 88
=====
Sap IP Address Configured MAC Dynamic MAC
Subscriber Fwding state

1/2/4:50/5 10.144.145.1 N/A N/A
N/A Fwding

Number of static hosts : 1
=====
*A:ALA-48#
```

**Table 22** Show All Service-ID Field Descriptions

| Label            | Description                                                                        |
|------------------|------------------------------------------------------------------------------------|
| Service Id       | The service identifier.                                                            |
| VPN Id           | The number which identifies the VPN.                                               |
| Service Type     | Specifies the type of service.                                                     |
| SDP Id           | The SDP identifier.                                                                |
| Description      | Generic information about the service.                                             |
| Customer Id      | The customer identifier.                                                           |
| Last Mgmt Change | The date and time of the most recent management-initiated change to this customer. |
| SAP Count        | The number of SAPs specified for this service.                                     |

## host-connectivity-verify

**Syntax** **host-connectivity-verify**

**Context** show>service>id

**Description** This command enables the context to display host connectivity check statistics.

## statistics

**Syntax** **statistics [summary] [sap sap-id]**

**Context** show>service>id>host-connectivity-verify

**Description** This command displays host connectivity verification statistics.

**Parameters** *sap-id* — Specifies the SAP ID to show statistics for.

**summary** — Displays host connectivity verify summary information.

**Output** The following output is an example of host connectivity verification statistics information, and [Table 23](#) describes the output fields.

### Sample Output

```
A:ALA-48>show>service>id# host-connectivity-verify statistics sap 1/1/9:0
=====
Host connectivity check statistics
=====
Svc SapId/ DestIp Last Time Oper
Id SdpId Address Response Expired State

1000 551/2/3:0 10.144.145.1 Up
=====
A:ALA-48>show>service>id#
```

**Table 23** Show Host Connectivity Verify Statistics Field Descriptions

| Label          | Description                                            |
|----------------|--------------------------------------------------------|
| Svc Id         | The service identifier.                                |
| SapId/SdpId    | The SAP and SDP identifiers.                           |
| DestIp Address | The destination IP address.                            |
| Last Response  | The time when the last response was received.          |
| Time Expired   | Displays whether the interval value has expired.       |
| Oper State     | Displays the current operational state of the service. |

## summary

- Syntax** **summary** [**sap** *sap-id*]
- Context** show>service>id>host-connectivity-verify
- Description** This command displays a host connectivity verification summary.
- Parameters** *sap-id* — Specifies the SAP ID to show statistics for.

## interface

- Syntax** **interface** [{*[ip-address | ip-int-name]* [*interface-type*] [*detail*] [*family*]}] | **summary**
- Context** show>service>id

**Description** This command displays information for the IP interfaces associated with the IES service. If no optional parameters are specified, a summary of all IP interfaces associated to the service are displayed.

**Parameters** *ip-address* — The IP address of the interface for which to display information.

**Values**

|              |              |                                     |
|--------------|--------------|-------------------------------------|
| <ip-address> | 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. 32 characters maximum.

**family** — Displays the router IP interface table to display.

**Values** **ipv4** — Displays only those peers that have the IPv4 family enabled.  
**ipv6** — Displays the peers that are IPv6-capable.

*interface-type* — Specifies whether to display groups, interfaces, or redundant interfaces.

**Values** group, subscriber, redundant

**detail** — Displays detailed IP interface information.

**summary** — Displays IP interface information summary.

**Output** The following output is an example of IP interface information, and [Table 24](#) describes the output fields.

**Sample Output**

```
A:ALA-49# show service id 88 interface
=====
Interface Table
=====
```

| Interface-Name<br>IP-Address | Adm | Opr (v4/v6) | Type    | Port/SapId<br>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 24 Show Service-ID Field Descriptions**

| Label           | Description                                                                                                                               |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| If Name         | The name used to refer to the IES interface.                                                                                              |
| Type            | Specifies the interface type.                                                                                                             |
| IP-Address      | Specifies the IP address/IP subnet/broadcast address of the interface.                                                                    |
| Adm             | The administrative state of the interface.                                                                                                |
| Opr             | The operational state of the interface.                                                                                                   |
| Admin State     | The administrative state of the interface.                                                                                                |
| Oper State      | The operational state of the interface.                                                                                                   |
| IP Addr/mask    | Specifies the IP address/IP subnet/broadcast address of the interface.                                                                    |
| If Index        | The index corresponding to this IES interface. The primary index is 1; all IES interfaces are defined in the base virtual router context. |
| If Type         | Specifies the interface type.                                                                                                             |
| SAP Id          | Specifies the SAP's port ID.                                                                                                              |
| SNTP B.Cast     | Specifies whether SNTP broadcast client mode is enabled or disabled.                                                                      |
| Arp Timeout     | Specifies the timeout for an ARP entry learned on the interface.                                                                          |
| 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.                                                                                                 |

labels

**Syntax labels**



- Context** show>service>id
- Description** Displays the labels being used by the service.
- Output** The following output is an example of service ID labels information, and [Table 25](#) describes the output fields.

### Sample Output

```
*A:ALA-12# show service id 1 labels
=====
Martini Service Labels
=====
Svc Id Sdp Id Type I.Lbl E.Lbl

1 10:1 Mesh 0 0
1 20:1 Mesh 0 0
1 30:1 Mesh 0 0
1 40:1 Mesh 130081 131061
1 60:1 Mesh 131019 131016
1 100:1 Mesh 0 0

Number of Bound SDPs : 6

*A:ALA-12#
```

**Table 25** Show Service-ID Labels Field Descriptions

| Label  | Description                                                                                        |
|--------|----------------------------------------------------------------------------------------------------|
| Svc Id | The service identifier.                                                                            |
| Sdp Id | The SDP identifier.                                                                                |
| Type   | Indicates whether the SDP is a spoke or a mesh.                                                    |
| I.Lbl  | The VC label used by the far-end device to send packets to this device in this service by the SDP. |
| E.Lbl  | The VC label used by this device to send packets to the far-end device in this service by the SDP. |

## sdp

- Syntax**
- sdp sdp-id pw-port [pw-port-id]**
  - sdp sdp-id pw-port**
  - sdp sdp-id pw-port [pw-port-id] [statistics]**
  - sdp [consistent | inconsistent | na] egressifs**
  - sdp sdp-id keep-alive-history**
  - sdp far-end {ip-address | ipv6-address} keep-alive-history**

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                    | <b>sdp [sdp-id] [detail]</b><br><b>sdp far-end {ip-address   ipv6-address} [detail]</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | show>service                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>Displays information for the SDPs associated with the service.</p> <p>If no optional parameters are specified, a summary of all associated SDPs is displayed.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <p><b>sdp-id</b> — Specifies the SDP ID for which to display information.</p> <p><b>Values</b> 1 to 17407</p> <p><b>Default</b> All SDPs.</p> <p><b>pw-port-id</b> — Specifies the pseudo-wire port identifier.</p> <p><b>Values</b> 1 to 10239</p> <p><b>statistics</b> — Displays SDP statistics information.</p> <p><b>consistent</b> — Indicates that the network-domains for all the egress network interfaces that can carry traffic on this SDP are consistent.</p> <p><b>inconsistent</b> — Indicates that the network-domain for one or more egress network interfaces that can carry traffic on this SDP are inconsistent.</p> <p><b>na</b> — Indicates that there is no egress network interface that can carry traffic on this SDP.</p> <p><b>egressifs</b> — Indicates whether all the egress network interfaces that can carry traffic on this SDP are associated with the network-domain configured on this SDP.</p> <p><b>ip-address   ipv6-address</b> — Displays only SDPs matching with the specified far-end IP address. 64 characters maximum.</p> <p><b>Default</b> SDPs with any far-end IP address.</p> <p><b>keep-alive-history</b> — Displays the last fifty SDP keepalive events for the SDP.</p> <p><b>Default</b> SDP summary output.</p> <p><b>detail</b> — Displays detailed SDP information.</p> <p><b>Default</b> SDP summary output.</p> |
| <b>Output</b>      | <p>The following output is an example of SDP information, and <a href="#">Table 17</a> describes the output fields.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

### Sample Output

```
*A:ALA-12>config>service# show service sdp 1 pw-port
=====
Service Destination Point (sdp Id 1 Pw-Port)
=====
Pw-port VC-Id Adm Encap Opr VC Type Egr Monitor
 Shaper Oper
 VPort Group

1 1 up dot1q up ether
```

```

2 2 up qinq up ether
3 3 up dot1q up ether
4 4 up qinq up ether

```

-----  
Entries found : 4  
=====

\*A:ALA-12>config>service# show service sdp 1 pw-port 3

=====

Service Destination Point (Sdp Id 1 Pw-Port 3)

```

=====
SDP Binding port : lag-1
VC-Id : 3 Admin Status : up
Encap : dot1q Oper Status : up
VC Type : ether
Oper Flags : (Not Specified)
Monitor Oper-Group : (Not Specified)
=====

```

\*A:ALA-12>config>service# show service sdp 1 pw-port 3 statistics

=====

Service Destination Point (Sdp Id 1 Pw-Port 3)

```

=====
SDP Binding port : lag-1
VC-Id : 3 Admin Status : up
Encap : dot1q Oper Status : up
VC Type : ether
Oper Flags : (Not Specified)
Monitor Oper-Group : (Not Specified)
=====

```

```

Statistics :
I. Fwd. Pkts. : 0 I. Dro. Pkts. : 0
I. Fwd. Octs. : 0 I. Dro. Octs. : 0
E. Fwd. Pkts. : 0 E. Fwd. Octets : 0
=====

```

\*A:Dut-B# show service sdp

=====

Services: Service Destination Points

```

=====
SdpId AdmMTU OprMTU Far End Adm Opr Del LSP Sig

230 0 1582 10.20.1.3 Up Up MPLS I TLDP
=====

```

Number of SDPs : 1

-----

Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable

=====

\*A:Dut-B#

\*A:Dut-B# show service sdp detail

=====

Services: Service Destination Points Details

=====

```

Sdp Id 230 -10.20.1.3

Description : (Not Specified)
SDP Id : 230 SDP Source : manual
Admin Path MTU : 0 Oper Path MTU : 1582
Delivery : MPLS
Far End : 10.20.1.3
Tunnel Far End : n/a LSP Types : SR-ISIS

Admin State : Up Oper State : Up
Signaling : TLDP Metric : 0
Acct. Pol : None Collect Stats : Disabled
Last Status Change : 01/28/2015 22:00:07 Adv. MTU Over. : No
Last Mgmt Change : 01/28/2015 21:59:53 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 : SR-ISIS

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

Src B-MAC LSB : <none> Ctrl PW VC ID : <none>
Ctrl PW Active : n/a

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

Segment Routing

ISIS : enabled LSP Id : 524289
Oper Instance Id : 0

Number of SDPs : 1

```

```

=====
*A:Dut-B#

*A:Dut-B>config>service>sdp# show service sdp

=====
Services: Service Destination Points
=====
SdpId AdmMTU OprMTU Far End Adm Opr Del LSP Sig

230 0 1582 10.20.1.3 Up Up MPLS O TLDP

Number of SDPs : 1

Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable
 I = SR-ISIS, O = SR-OSPF
=====

*A:ALA-12# show service sdp 2 detail
=====
Service Destination Point (Sdp Id : 2) Details

Sdp Id 2 -(10.10.10.104)

Description : GRE-10.10.10.104
SDP Id : 2
Admin Path MTU : 0 Oper Path MTU : 0
Far End : 10.10.10.104 Delivery : GRE
Admin State : Up Oper State : Down
Weighted ECMP : Disabled VLAN VC Etype : 0x8100
Flags : SignalingSessDown TransportTunnDown
Signaling : TLDP Adv. MTU Over. : No
Last Status Change : 02/01/2007 09:11:39
Last Mgmt Change : 02/01/2007 09:11:46

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

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
=====
*A:ALA-12#

```

**Table 26** Show Service SDP Field Descriptions

| Label  | Description         |
|--------|---------------------|
| SDP Id | The SDP identifier. |

**Table 26 Show Service SDP Field Descriptions (Continued)**

| Label                 | Description (Continued)                                                                                                                                                     |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Adm MTU               | Specifies the desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented. |
| 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 GRE or MPLS tunnel defined by this SDP.                                                                                   |
| Adm                   | Admin State — Specifies the administrative state of the SDP.                                                                                                                |
| Opr                   | Oper State — Specifies the operational state of the SDP.                                                                                                                    |
| Deliver               | Specifies the type of delivery used by the SDP: GRE or MPLS.                                                                                                                |
| Flags                 | Specifies the conditions that affect the operating status of this SDP.                                                                                                      |
| Signal<br>Signaling   | Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.                                           |
| Last Status<br>Change | Specifies the time of the most recent operating status change to this SDP.                                                                                                  |

## pw-port

**Syntax** **pw-port** [*pw-port-id*] [**detail**]  
**pw-port sdp** *sdp-id*  
**pw-port sdp none**  
**pw-port** *pw-port-id* [**statistics**]

**Context** show>pw-port

**Description** Displays pseudo-wire port information.

If no optional parameters are specified, the command displays a summary of all defined PW ports. The optional parameters restrict output to only ports matching the specified properties.

**Parameters** *pw-port-id* — Specifies the pseudo-wire port identifier.

**Values** 1 to 10239

**detail** — Displays detailed port information that includes all the **pw-port** output fields.

*sdp-id* — The SDP ID for which to display matching PW port information.

**Values** 1 to 17407

**statistics** — Displays statistics information.

**Output** The following output is an example of PW port information, and [Table 27](#) describes the output fields.

### Sample Output

```
*A:ALA-48>config>service# show pw-port
=====
PW Port Information
=====
PW Port Encap SDP IfIndex VC-Id

1 dot1q 1 1526726657 1
2 qinq 1 1526726658 2
3 dot1q 1 1526726659 3
4 qinq 1 1526726660 4
=====

*A:ALA-48>config>service# show pw-port 3
=====
PW Port Information
=====
PW Port Encap SDP IfIndex VC-Id

3 dot1q 1 1526726659 3
=====

*A:ALA-48>config>service# show pw-port 3 detail
=====
PW Port Information
=====
PW Port : 3
Encap : dot1q
SDP : 1
IfIndex : 1526726659
VC-Id : 3
Description : 1-Gig Ethernet dual fiber
=====

*A:ALA-48>config>pw-port$ show pw-port sdp none
=====
PW Port Information
=====
PW Port Encap SDP IfIndex VC-Id

5 dot1q 1526726661
=====

*A:ALA-48>config>pw-port$ show pw-port sdp 1
=====
PW Port Information
=====
PW Port Encap SDP IfIndex VC-Id

1 dot1q 1 1526726657 1
```

```

2 qinq 1 1526726658 2
3 dot1q 1 1526726659 3
4 qinq 1 1526726660 4
=====

```

**Table 27** Show PW-Port Field Descriptions

| Label       | Description                               |
|-------------|-------------------------------------------|
| PW Port     | The PW port identifier.                   |
| Encap       | The encapsulation type of the PW port.    |
| SDP         | The SDP identifier.                       |
| IfIndex     | The interface index used for the PW port. |
| VC-Id       | The virtual circuit identifier.           |
| Description | The description string for the PW port.   |

## sap

- Syntax** `sap sap-id`
- Context** `show>service>id`
- Description** This command displays information for the SAPs associated with the service.
- If no optional parameters are specified, a summary of all associated SAPs is displayed.
- Parameters** *sap-id* — The ID that displays SAPs for the service in the *slot/mdal/port[channel]* format.
- Output** The following output is an example of SAP information, and [Table 28](#) describes the output fields.

### Sample Output

```

A:PE# show service id 500 sap 1/1/9:1 detail

=====
Service Access Points(SAP)
=====
Service Id : 500
SAP : 1/1/9:1 Encap : q-tag
Description : (Not Specified)
Admin State : Up Oper State : Up
Flags : None
Multi Svc Site : None
Last Status Change : 08/31/2018 16:07:55
Last Mgmt Change : 08/31/2018 16:07:55
Sub Type : regular

```



|                                             |                                  |
|---------------------------------------------|----------------------------------|
| Dot1Q Ethertype : 0x8100                    | QinQ Ethertype : 0x8100          |
| Split Horizon Group: (Not Specified)        |                                  |
| Admin MTU : 1518                            | Oper MTU : 1518                  |
| Ingr IP Fltr-Id : n/a                       | Egr IP Fltr-Id : n/a             |
| Ingr Mac Fltr-Id : n/a                      | Egr Mac Fltr-Id : n/a            |
| Ingr IPv6 Fltr-Id : n/a                     | Egr IPv6 Fltr-Id : n/a           |
| qinq-pbit-marking : both                    |                                  |
| Egr Agg Rate Limit : max                    |                                  |
| Q Frame-Based Acct : Disabled               | Limit Unused BW : Disabled       |
| Qinq-vlan-translation : None                | Qinq-vlan-translation Ids : None |
| Acct. Pol : None                            | Collect Stats : Disabled         |
| Anti Spoofing : None                        | Dynamic Hosts : Enabled          |
| Avl Static Hosts : 0                        | Tot Static Hosts : 0             |
| Calling-Station-Id : n/a                    |                                  |
| Application Profile: None                   |                                  |
| Transit Policy : None                       |                                  |
| AARP Id : None                              |                                  |
| Oper Group : (none)                         | Monitor Oper Grp : (none)        |
| Host Lockout Plcy : n/a                     |                                  |
| Lag Link Map Prof : (none)                  |                                  |
| Bandwidth : Not-Applicable                  |                                  |
| Oper DCpu Prot Pol*: _default-access-policy |                                  |

-----  
ETH-CFM SAP specifics  
-----

|                              |                |
|------------------------------|----------------|
| Tunnel Faults : n/a          | AIS : Disabled |
| MC Prop-Hold-Timer : n/a     |                |
| Squelch Levels : None        |                |
| Collect Lmm Stats : Disabled |                |
| LMM FC Stats : None          |                |
| LMM FC In Prof : None        |                |

-----  
QOS  
-----

|                                      |                              |
|--------------------------------------|------------------------------|
| Ingress qos-policy : 1               | Egress qos-policy : 1        |
| Ingress FP QGrp : (none)             | Egress Port QGrp : (none)    |
| Ing FP QGrp Inst : (none)            | Egr Port QGrp Inst: (none)   |
| Shared Q plcy : n/a                  | Multipoint shared : Disabled |
| I. Sched Pol : (Not Specified)       |                              |
| E. Sched Pol : (Not Specified)       |                              |
| I. Policer Ctl Pol : (Not Specified) |                              |
| E. Policer Ctl Pol : (Not Specified) |                              |
| E. HS Sec. Shaper : (Not Specified)  |                              |
| I. QGrp Redir. List: (Not Specified) |                              |
| E. QGrp Redir. List: (Not Specified) |                              |

-----  
Sap Aggregate Stats  
-----

|         | Packets | Octets |
|---------|---------|--------|
| Ingress |         |        |

|                     |     |   |
|---------------------|-----|---|
| Aggregate Offered   | : 0 | 0 |
| Aggregate Forwarded | : 0 | 0 |
| Aggregate Dropped   | : 0 | 0 |

|                     |     |   |
|---------------------|-----|---|
| Egress              |     |   |
| Aggregate Forwarded | : 0 | 0 |
| Aggregate Dropped   | : 0 | 0 |

-----  
Sap Statistics  
-----

Last Cleared Time : 08/31/2018 16:19:29

|             |         |        |
|-------------|---------|--------|
|             | Packets | Octets |
| CPM Ingress | : 0     | 0      |

## Forwarding Engine Stats

|                |     |   |
|----------------|-----|---|
| Dropped        | : 0 | 0 |
| Received Valid | : 0 | 0 |
| Off. HiPrio    | : 0 | 0 |
| Off. LowPrio   | : 0 | 0 |
| Off. Uncolor   | : 0 | 0 |
| Off. Managed   | : 0 | 0 |

## Queueing Stats(Ingress QoS Policy 1)

|              |     |   |
|--------------|-----|---|
| Dro. HiPrio  | : 0 | 0 |
| Dro. LowPrio | : 0 | 0 |
| For. InProf  | : 0 | 0 |
| For. OutProf | : 0 | 0 |

## Queueing Stats(Egress QoS Policy 1)

|                    |     |   |
|--------------------|-----|---|
| Dro. In/InplusProf | : 0 | 0 |
| Dro. Out/ExcProf   | : 0 | 0 |
| For. In/InplusProf | : 0 | 0 |
| For. Out/ExcProf   | : 0 | 0 |

-----  
Sap per Queue stats  
-----

|                                      |         |        |
|--------------------------------------|---------|--------|
|                                      | Packets | Octets |
| Ingress Queue 1 (Unicast) (Priority) |         |        |
| Off. HiPrio                          | : 0     | 0      |
| Off. LowPrio                         | : 0     | 0      |
| Dro. HiPrio                          | : 0     | 0      |
| Dro. LowPrio                         | : 0     | 0      |
| For. InProf                          | : 0     | 0      |
| For. OutProf                         | : 0     | 0      |

## Egress Queue 1

|                    |     |   |
|--------------------|-----|---|
| For. In/InplusProf | : 0 | 0 |
| For. Out/ExcProf   | : 0 | 0 |
| Dro. In/InplusProf | : 0 | 0 |
| Dro. Out/ExcProf   | : 0 | 0 |

=====

\* indicates that the corresponding row element may have been truncated.

A:PE#

**Table 28 Show Service-ID SAP Field Descriptions**

| Label              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Service Id         | The service identifier.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| SAP                | The type of SAP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Encap              | The encapsulation type of the SAP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Ethertype          | Specifies an Ethernet type II Ethertype value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Admin State        | The administrative state of the SAP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Oper State         | The operational state of the SAP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Flags              | Specifies the conditions that affect the operating status of this SAP.<br>Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapPipeCelpAddr, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode, SapIngQGrpRedirMismatch, SapEgrQGrpRedirMismatch |
| 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.                                                                                                                                                                                                                                                                                                                                                                      |
| 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 statistics collection is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Dropped            | The number of packets and octets dropped due to SAP state, ingress MAC or IP filter, same segment discard, bad checksum, and so on.                                                                                                                                                                                                                                                                                                                                                                                                   |
| Received Valid     | The number of valid packets and octets received on the SAP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Off. HiPrio        | The number of high-priority packets and octets, as determined by the SAP ingress QoS policy.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Off. LowPrio       | The number of low-priority packets and octets, as determined by the SAP ingress QoS policy.                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Table 28** Show Service-ID SAP Field Descriptions (Continued)

| Label               | Description (Continued)                                                                                                                                                       |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Off. Uncolor        | The number of uncolored packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.                                                   |
| Dro. HiPrio         | The number of high-priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to MBS exceeded, buffer pool limit exceeded, and so on. |
| Dro. LowPrio        | The number of low-priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to MBS exceeded, buffer pool limit exceeded, and so on.  |
| For. InProf         | The number of in-profile packets and octets (rate below CIR) forwarded by the ingress Qchip.                                                                                  |
| For. OutProf        | The number of out-of-profile packets and octets (rate below CIR) forwarded by the ingress Qchip.                                                                              |
| 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.                                         |
| 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 OAM operational status of the VCL.                                                                                                                                        |
| AAL-5 Encap         | The AAL-5 encapsulation type.                                                                                                                                                 |
| Mult Svc Site       | Specifies the customer's multiservice site name.                                                                                                                              |
| I. Sched Pol        | The ingress scheduler policy applied to the customer's multiservice site.                                                                                                     |
| E. Sched Pol        | The egress scheduler policy applied to the customer's multiservice site.                                                                                                      |

## sap

**Syntax** `sap sap-id queue-depth [queue queue-id] [ingress | egress]`

**Context** `show>service>id`

**Description** This command displays queue depth information for the specified SAP.

**Parameters** *sap-id* — The ID that displays SAPs for the service in the *slot/mdalport[.channel]* format.  
*queue-id* — Specifies the queue ID.

**Values** 1 to 32

**ingress** — Displays information for the ingress policy.

**egress** — Displays information for the egress policy.

**Output** The following output is an example of SAP queue depth information.

### Sample Output

```
*A:PE-1# show service id 1 sap 1/2/1 queue-depth
=====
Queue Depth Information (Ingress SAP)
=====
No Matching Entries
=====
Queue Depth Information (Egress SAP)
=====

Name : 1->1/2/1->1
MBS : Def

Queue Depths (percentage)

0%-10% 11%-20% 21%-30% 31%-40% 41%-50% 51%-60% 61%-70% 71%-80% 81%-90% 91%-100%

68.21 3.64 3.43 3.47 3.86 3.22 3.86 2.87 3.78 3.66

Average Elapsed Time : 0d 00:11:48
Wghtd Avg Polling Interval: 99 ms

=====
*A:PE-1#
```

## sap

- Syntax** **sap** *sap-id* **queue-group-redirection** [**ingress** | **egress**]
- Context** show>service>id
- Description** This command displays queue group redirect information for the specified SAP.
- The output lists the queue group name and the instances configured in the related queue group redirect list. For each instance, the FP (for ingress) and port (for egress) is displayed. If there is a mismatch between the SAP and redirect list configuration and the queue group instance configuration, this is highlighted.
- Parameters** *sap-id* — The ID that displays SAPs for the service in the *slot/mdalport[.channel]* format.
- ingress** — Displays information for the ingress policy.
- egress** — Displays information for the egress policy.
- Output** The following output is an example of SAP queue group redirection information.

**Sample Output**

```

*A:PE# show service id 1 sap 1/1/1 queue-group-redirect
=====
Queue Group Redirect List Information (Ingress SAP)
=====
Queue Group Redirect List : list1
Type : vxlan-vni
Queue Group : qg1

Match Instance FP

1 1 1/1
2 2 1/1
3 3 1/1 : mismatch

=====
Queue Group Redirect List Information (Egress SAP)
=====
Queue Group Redirect List : list1
Type : vxlan-vni
Queue Group : qg1

Match Instance Port

1 1 1/1/1
2 2 1/1/1
3 3 1/1/1 : mismatch

*A:PE#

```

**sdp**

**Syntax**

```

sdp [sdp-id [:vc-id]] [detail]
sdp far-end {ip-address | ipv6-address} [detail]
sdp sdp-id [vc-id] l2tpv3
sdp sdp-id [:vc-id] static-isids [range-id range-id]
sdp sdp-id [:vc-id] static-isids mfib
sdp sdp-id [:vc-id] [detail] vccv-bfd [session]
sdp sdp-id [:vc-id] mrp

```

**Context** show>service>id

**Description** This command displays information for the SDPs associated with the service.

If no optional parameters are specified, a summary of all associated SDPs is displayed.

- Parameters**    *sdp-id* [:*vc-id*] — Displays only information for the specified SDP ID.
- Values**        sdp-id: 1 to 17407  
                                      vc-id: 1 to 4294967295
- Default**      all SDPs
- ip-address* | *ipv6-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.
- l2tpv3** — Indicates that the user wants to display l2tpv3 specific information for SDPs that are of type l2tpv3.
- static-isids** — Specifies the I-Component service IDs created on the SDP.
- range-id* — Displays the service using the specified I-component Service ID (ISID)
- Values**        1 to 4294967295
- mfib** — Displays MFIB related information. This parameter applies to the 7450 ESS or 7750 SR only.
- vccv-bfd** — Displays detailed information about the VCCV BFD session for a spoke SDP.
- session** — Displays a summary of all VCCV sessions.
- mrp** — Displays detailed MRP information.

**Output**        The following output is an example of SDP information.

### 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
Learned MAC Addr : 0
Total MAC Addr : 0
Static MAC Addr : 0

MAC Learning : Enabled
MAC Aging : Enabled
L2PT Termination : Disabled
MAC Pinning : Disabled
Discard Unkwn Srce : Disabled
BPDU Translation : Disabled

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

Statistics :
I. Fwd. Pkts. : 0
I. Fwd. Octs. : 0
E. Fwd. Pkts. : 0
E. Fwd. Octets : 0
MCAC Policy Name :
MCAC Max Unconst BW : no limit
MCAC In use Mand BW : 0
MCAC In use Opnl BW : 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 Oper State : Up
Time Since Last Tr*: 00h26m34s

Lsp Name : A_B_10
Admin State : Up Oper State : Up
Time Since Last Tr*: 00h26m34s

Class-based forwarding :

Class forwarding : enabled
Default LSP : A_B_10 Multicast LSP : A_B_9
=====
FC Mapping Table
=====
FC Name LSP Name

af A_B_3
be A_B_1
ef A_B_6
h1 A_B_7
h2 A_B_5
l1 A_B_4
l2 A_B_2
nc A_B_8
=====
Stp Service Destination Point specifics

Mac Move : Blockable
Stp Admin State : Up Stp Oper State : Down
Core Connectivity : Down
Port Role : N/A Port State : Forwarding
Port Number : 2049 Port Priority : 128
Port Path Cost : 10 Auto Edge : Enabled
Admin Edge : Disabled Oper Edge : N/A
Link Type : Pt-pt BPDU Encap : Dot1d
Root Guard : Disabled Active Protocol : N/A
Last BPDU from : N/A
Designated Bridge : N/A Designated Port Id: 0

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

Number of SDPs : 1

* indicates that the corresponding row element may have been truncated.

A:Dut-A#

```

## subscriber-hosts

**Syntax**     **subscriber-hosts** [sap sap-id] [ip ip-prefix/prefix-length] [mac ieee-address] [sub-profile sub-profile-name] [sla-profile sla-profile-name] [app-profile app-profile-name]

---

**[wholesaler service-id] [address-origin address-origin] [detail]**

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | show>service>id                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | This command displays subscriber host information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <p><i>sap-id</i> — Displays the specified subscriber host SAP information.</p> <p><i>ip-prefix/prefix-length</i> — The destination address of the aggregate route in dotted decimal notation.</p> <p><b>Values</b>      ipv4-prefix: a.b.c.d (host bits must be 0)<br/>                           ipv4-prefix-length: 0 to 32<br/>                           ipv6-prefix:<br/>                             • x:x:x:x:x:x:x (eight 16-bit pieces)<br/>                             • x:x:x:x:x:x.d.d.d.d<br/>                             • x: [0 to FFFF] H<br/>                             • d: [0 to 255] D<br/>                           ipv6-prefix-length: 0 to 128</p> <p><i>ieee-address</i> — 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.</p> <p><b>Values</b>      xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx (cannot be all zeros)</p> <p><i>sub-profile-name</i> — Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the <b>config&gt;subscr-mgmt&gt;sub-profile</b> context. 32 characters maximum.</p> <p><i>sla-profile-name</i> — Specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the <b>config&gt;subscr-mgmt&gt;sla-profile</b> context. 32 characters maximum.</p> <p><i>app-profile-name</i> — Specifies an existing application profile name to be associated with the static subscriber host. The application profile is configured in the <b>config&gt;subscr-mgmt&gt;sla-profile</b> context. 32 characters maximum.</p> <p><i>service-id</i> — The VPRN service ID of the wholesaler (applies only to the 7750 SR).</p> <p><b>Values</b>      1 to 2147483648</p> <p><i>address-origin</i> — Shows the origin of the IP address assigned to the subscriber-host.</p> <p><b>Values</b>      aaa, dynamic, static, bonding</p> <p><b>detail</b> — Displays detailed information.</p> |

## statistics

**Syntax**      **statistics** [*ip-int-name* | *ip-address*]

**Context** show>router>dhcp

**Description** Display statistics for DHCP relay and DHCP snooping. If no IP address or interface name is specified, then all configured interfaces are displayed. If an IP address or interface name is specified, then only data regarding the specified interface is displayed.

**Parameters** *ip-int-name* | *ip-address* — Displays statistics for the specified IP interface.

**Output** The following output is an example of DHCP statistics information, and [Table 29](#) describes the output fields.

### Sample Output

```
*A:ALA-1# show router dhcp statistics
=====
DHCP Global Statistics
=====
Rx Packets : 0
Tx Packets : 0
Rx Malformed Packets : 0
Rx Untrusted Packets : 0
Client Packets Discarded : 0
Client Packets Relayed : 0
Client Packets Snooped : 0
Server Packets Discarded : 0
Server Packets Relayed : 0
Server Packets Snooped : 0
=====
*A:ALA-1#
```

**Table 29 Show DHCP Statistics Field Descriptions**

| Label                      | Description                                                               |
|----------------------------|---------------------------------------------------------------------------|
| Received Packets           | The number of packets received from the DHCP clients.                     |
| Transmitted Packets        | The number of packets transmitted to the DHCP clients.                    |
| Received Malformed Packets | The number of malformed packets received from the DHCP clients.           |
| Received Untrusted Packets | The number of untrusted packets received from the DHCP clients.           |
| 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

|                    |                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>summary</b>                                                                                                            |
| <b>Context</b>     | show>router>dhcp                                                                                                          |
| <b>Description</b> | This command displays the status of the DHCP relay and DHCP snooping functions on each interface.                         |
| <b>Output</b>      | The following output is an example of DHCP summary information, and <a href="#">Table 30</a> describes the output fields. |

**Sample Output**

The following sample output applies to the 7750 SR:

```
A:ALA-49# show router dhcp summary
=====
DHCP Summary (Router: Base)
=====
```

| Interface Name<br>SapId/Sdp | Arp<br>Populate | Used/<br>Provided | Info<br>Option | Admin<br>State |
|-----------------------------|-----------------|-------------------|----------------|----------------|
| Sector A                    | No              | 0/0               | Keep           | Up             |
| sap:7/1/1.2.2               |                 | 0/0               |                |                |
| grp-if                      | No              | 0/1               | Keep           | Down           |
| ies-test                    | No              | 0/0               | Keep           | Up             |
| sap:9/1/2:0/500             |                 | 0/0               |                |                |
| test                        | No              | 0/0               | Keep           | Up             |
| sap:10/1/2:0                |                 | 0/0               |                |                |
| test1                       | No              | 0/0               | Keep           | Up             |
| sap:7/1/1.1.2               |                 | 0/0               |                |                |
| test2                       | No              | 0/0               | Keep           | Up             |
| sap:7/1/1.2.1               |                 | 0/0               |                |                |
| testA                       | No              | 0/0               | Keep           | Up             |
| sap:7/1/3.1.1               |                 | 0/0               |                |                |
| testB                       | No              | 0/0               | Keep           | Up             |
| sap:7/1/5.1.1               |                 | 0/0               |                |                |
| to-HQ                       | No              | 0/0               | Keep           | Up             |
| sdp:spoke-2:1001            |                 | 0/0               |                |                |
| to-web                      | No              | 0/0               | Keep           | Up             |
| sap:2/1/10:50               |                 | 0/0               |                |                |

```

Interfaces: 9
```

**Sample Output**

The following sample output applies to the 7450 ESS:

```
A:ALA-49# show router dhcp summary
=====
DHCP Summary (Router: Base)
=====
```

| Interface Name<br>SapId/Sdp | Arp<br>Populate | Used/<br>Provided | Info<br>Option | Admin<br>State |
|-----------------------------|-----------------|-------------------|----------------|----------------|
| Sector A                    | No              | 0/0               | Keep           | Up             |

```

 sap:7/1/1.2.2 0/0
ies-test No 0/0 Keep Up
 sap:9/1/2:0/500 0/0
test No 0/0 Keep Up
 sap:10/1/2:0 0/0
test1 No 0/0 Keep Up
 sap:7/1/1.1.2 0/0
test2 No 0/0 Keep Up
 sap:7/1/1.2.1 0/0
testA No 0/0 Keep Up
 sap:7/1/3.1.1 0/0
to-HQ No 0/0 Keep Up
 sdp:spoke-2:1001 0/0
to-web No 0/0 Keep Up
sap:2/1/10:50 0/0

Interfaces: 9

```

**Table 30** Show DHCP Summary Field Descriptions

| Label          | Description                                                                                                                                                                                                                                                             |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface Name | Name of the router interface.                                                                                                                                                                                                                                           |
| Sapld/Sdp      | Specifies the associated SAP ID or SDP ID.                                                                                                                                                                                                                              |
| Arp Populate   | Specifies whether or not ARP populate is enabled.                                                                                                                                                                                                                       |
| Used/Provided  | Used — The number of lease-states currently in use on a specific interface (the number of clients on that interface got an IP address by DHCP). This value is always less than or equal to the 'Provided' field.<br>Provided — The configured for a specific interface. |
| Info Option    | Indicates whether Option 82 processing is enabled on the interface.                                                                                                                                                                                                     |
| Admin State    | Indicates the administrative state.                                                                                                                                                                                                                                     |

## vrrp

**Syntax** **vrrp**

**Context** show>router

**Description** This command displays information VRRP instances.

## instance

**Syntax** **instance**

**instance interface** *interface-name* [**vrid** *virtual-router-id*]

**instance interface** *interface-name* **vrid** *virtual-router-id* **ipv6**

---

|                    |                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | show>router>vrrp                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command displays statistics for the VRRP instance.                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <i>interface-name</i> — Displays statistics for the specified interface, up to 32 characters maximum.<br><i>virtual-router-id</i> — Displays statistics for the specified virtual router ID.<br><b>Values</b> 1 to 255<br><b>ipv6</b> — Displays statistics peers for IPv6 instances. |

## statistics

|                    |                                                         |
|--------------------|---------------------------------------------------------|
| <b>Syntax</b>      | <b>statistics</b>                                       |
| <b>Context</b>     | show>router>vrrp                                        |
| <b>Description</b> | This command displays statistics for the VRRP instance. |

## retailers

|                    |                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>retailers</b>                                                                                          |
| <b>Context</b>     | show>service>id                                                                                           |
| <b>Description</b> | This command displays the service ID of the retailer subscriber service to which this DHCP lease belongs. |

## wholesalers

|                    |                                                       |
|--------------------|-------------------------------------------------------|
| <b>Syntax</b>      | <b>wholesalers</b>                                    |
| <b>Context</b>     | show>service>id                                       |
| <b>Description</b> | This command displays service wholesaler information. |

### 2.6.2.2 IES Clear Commands

## dhcp

|                |                   |
|----------------|-------------------|
| <b>Syntax</b>  | <b>dhcp</b>       |
| <b>Context</b> | clear>router>dhcp |

clear>service>id

**Description** This command enables the context to clear DHCP parameters.

## dhcp6

**Syntax** **dhcp6**

**Context** clear>router>dhcp6  
clear>service>id

**Description** This command enables the context to clear DHCPv6 parameters and only applies to the 7750 SR.

## statistics

**Syntax** *statistics [ip-int-name | ip-address]*

**Context** clear>router>dhcp

**Description** Clears DHCP statistics.

## id

**Syntax** **id service-id**

**Context** clear>service  
clear>service>statistics

**Description** This command clears parameters for a specific service.

**Parameters** *service-id* — The ID that uniquely identifies the service to clear.

**Values** 1 to 2148278316, *svc-name: 64 characters max*

## arp-host

**Syntax** **arp-host {all | mac *ieee-address* | sap *sap-id* | ip-address *ip-address*[/*mask*]}**  
**arp-host {port *port-id* | {inter-dest-id *intermediate-destination-id* | no-inter-dest-id} [port *port-id*]}**  
**arp-host statistics [sap *sap-id* | interface *interface-name*]**

**Context** clear>service>id

**Description** This command clears ARP host data.

---

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <p><i>ieee-address</i> — Clears the ARP host MAC address information. 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.</p> <p><i>sap-id</i> — Clears the specified SAP information.</p> <p><i>ip-address[/mask]</i> — Clears the specified IP address and mask.</p> <p><b>Values</b>      a.b.c.d.<br/>                  mask: 1 to 32</p> <p><i>port-id</i> — Clear the specified port ID information.</p> <p><i>intermediate-destination-id</i> — Displays information about the specified intermediate destination ID. 32 characters maximum.</p> <p><b>no-inter-dest-id</b> — Displays the information about no intermediate destination ID.</p> <p><i>interface-name</i> — Clears the interface name. 32 characters maximum.</p> |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## ip

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ip <i>ip-address</i></b>                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | debug>service>id>arp-host                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | This command displays ARP host events for a particular IP address.                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><i>ip-address</i> — The IP address of the IP interface. The <i>ip-address</i> portion of the <b>address</b> command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><b>Values</b>      a.b.c.d</p> |

## mac

|                    |                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] mac <i>ieee-address</i></b>                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | debug>service>id>arp-host                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | This command displays ARP host events for a particular MAC address.                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <p><i>mac-address</i> — 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.</p> <p><b>Values</b>      xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx (cannot be all zeros)</p> |



## mode

|                    |                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mode</b> { <b>all</b>   <b>dropped-only</b> }<br><b>no mode</b>                               |
| <b>Context</b>     | debug>service>id>arp-host                                                                        |
| <b>Description</b> | This command configures the ARP host tracing mode.                                               |
| <b>Parameters</b>  | <b>all</b> — Debugs all dropped packets.<br><b>dropped-only</b> — Only displays dropped packets. |

## sap

|                    |                                                                                       |
|--------------------|---------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [ <b>no</b> ] <b>sap</b> <i>sap-id</i>                                                |
| <b>Context</b>     | debug>service>id>host-connectivity-verify                                             |
| <b>Description</b> | This command displays ARP host events for a particular SAP.                           |
| <b>Parameters</b>  | <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. |

## interface

|                    |                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>interface</b> [ <i>ip-int-name</i>   <i>ip-addr</i> ] [ <b>icmp</b> ] [ <b>urpf-stats</b> ] [ <b>statistics</b> ]                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | clear>router                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command clears IP interface statistics.<br><br>If no IP interface is specified either by IP interface name or IP address, the command will perform the clear operation on all IP interfaces.                                                                                                                                                                                                 |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ip-addr</i> — The IP interface name or IP interface address.<br><br><b>Default</b> all IP interfaces<br><br><b>icmp</b> — Specifies to reset the ICMP statistics for the IP interface(s) used for ICMP rate limit.<br><br><b>urpf-stats</b> — Resets the statistics associated with uRPF failures.<br><br><b>statistics</b> — Resets the IP interface traffic statistics. |

## interface

|               |                                                                                                                                                                            |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>interface</b> <i>interface-name</i> [ <b>vrid</b> <i>virtual-router-id</i> ]<br><b>interface</b> <i>interface-name</i> <b>vrid</b> <i>virtual-router-id</i> <b>ipv6</b> |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

---

|                    |                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | clear>router>vrrp                                                                                                                                                  |
| <b>Description</b> | This command clears and resets VRRP instances.                                                                                                                     |
| <b>Parameters</b>  | <i>interface-name</i> — Specifies an existing interface name up to 32 characters in length.<br><i>virtual-router-id</i> — Specifies the virtual router identifier. |
| <b>Values</b>      | 1 to 255                                                                                                                                                           |
|                    | <b>ipv6</b> — Clears IPv6 VRRP interface information.                                                                                                              |

## statistics

|                    |                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>statistics interface</b> <i>interface-name</i> [ <b>vrid</b> <i>virtual-router-id</i> ]<br><b>statistics</b><br><b>statistics interface</b> <i>interface-name</i> <b>vrid</b> <i>virtual-router-id</i> <b>ipv6</b> |
| <b>Context</b>     | clear>router>vrrp                                                                                                                                                                                                     |
| <b>Description</b> | This command clears statistics for VRRP instances.                                                                                                                                                                    |
| <b>Parameters</b>  | <i>interface-name</i> — Specifies an existing interface name up to 32 characters in length.<br><i>virtual-router-id</i> — Specifies the virtual router identifier.                                                    |
| <b>Values</b>      | 1 to 255                                                                                                                                                                                                              |

## fdb

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fdb</b> { <b>all</b>   <b>mac</b> <i>ieee-address</i>   <b>sap</b> <i>sap-id</i> }   <b>mesh-sdp</b> <i>sdp-id</i> [: <i>vc-id</i> ]   <b>spoke-sdp</b> <i>sdp-id</i> [: <i>vc-id</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | clear>service>id                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | This command clears FDB entries for the service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b>  | <b>all</b> — Clears all FDB entries.<br><br><i>ieee-address</i> — Clears only FDB entries in the FDB table with the specified 48-bit MAC address. The MAC address can be expressed in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> where <i>aa</i> , <i>bb</i> , <i>cc</i> , <i>dd</i> , <i>ee</i> and <i>ff</i> are hexadecimal numbers.<br><br><i>sap-id</i> — Clears the specified SAP information.<br><br><b>mesh-sdp</b> — Clears only service FDB entries associated with the specified mesh SDP ID. For a mesh SDP, the VC ID is optional.<br><br><b>spoke-sdp</b> — Clears only service FDB entries associated with the specified spoke SDP ID. For a spoke SDP, the VC ID must be specified.<br><br><i>sdp-id</i> — The SDP ID for which to clear associated FDB entries. |
| <b>Values</b>      | 1 to 17407                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

*vc-id* — The virtual circuit ID on the SDP ID for which to clear associated FDB entries.

**Values** 1 to 4294967295

**Default** For mesh SDPs only, all VC IDs.

## site

**Syntax** **site** *name*

**Context** clear>service>id

**Description** This command clears site-specific information for the service.

**Parameters** *name* — Clears information about the specified service name. 32 characters maximum.

**Values** 1 to 2147483648

## spoke-sdp

**Syntax** **spoke-sdp** *sdp-id:vc-id ingress-vc-label*

**Context** clear>service>id

**Description** Clears and resets the spoke SDP bindings for the service.

**Parameters** *sdp-id* — The spoke SDP ID to be reset.

**Values** 1 to 17407

*vc-id* — The virtual circuit ID on the SDP ID to be reset.

**Values** 1 to 4294967295

## stp

**Syntax** **stp**

**Context** clear>service>statistics>id

**Description** Clears all spanning tree statistics for the service ID.

## lease-state

**Syntax** **lease-state ip-address** *ip-address[/mask]* [no-dhcp-release]

**lease-state mac** *ieee-address* [no-dhcp-release]

**lease-state mac sap** *sap-id* [no-dhcp-release]

**sdp sdp-id:vc-id [no-dhcp-release]**

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | clear>service>id>dhcp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | Clears DHCP lease state information for this service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><i>ip-address</i> [/mask] — Displays the specified destination IP address and mask.</p> <p><b>Values</b>      a.b.c.d<br/>                  mask: 1 to 32</p> <p><i>ieee-address</i> — 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.</p> <p><b>Values</b>      xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx</p> <p><i>sap-id</i> — Clears the specified lease state SAP information.</p> <p><i>sdp-id:vc-id</i> — The specified SDP to be cleared.</p> <p><b>Values</b>      sdp-id: 1 to 17407<br/>                  vc-id: 1 to 4294967295</p> <p><b>no-dhcp-release</b> — Specifies that the node will clear the state without sending the DHCP release message.</p> |

## lease-state

|                    |                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lease-state all [no-dhcp-release]</b><br><b>lease-state ipv6-address</b> <i>ipv6-prefix</i> [/prefix-length] [no-dhcp-release]<br><b>lease-state mac</b> <i>ieee-address</i> [no-dhcp-release]<br><b>lease-state sap</b> <i>sap-id</i> [no-dhcp-release] |
| <b>Context</b>     | clear>service>id>dhcp6                                                                                                                                                                                                                                      |
| <b>Description</b> | This command clears DHCPv6 lease state information for this service.                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><b>all</b> — Clears all statistics.</p> <p><b>no-dhcp-release</b> — Specifies that the node will clear the state without sending the DHCP release message.</p>                                                                                           |

*ipv6-prefix/prefix-length* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

**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

*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.

**Values**

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

## statistics

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>statistics</b> [ <i>ip-int-name</i>   <i>ipv6-address</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | clear>router>dhcp6                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | This command clears DHCP6 statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <p><i>ip-int-name</i> — Specifies the IP interface name up to 32 characters in length.</p> <p><i>ip-address ipv6-address/prefix-length</i> — The IP address of the IP interface. The <i>ip-address</i> portion of the <b>address</b> command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> |
| <b>Values</b>      | <p>ipv6-address: x:x:x:x:x:x:x (eight 16-bit pieces)<br/>x:x:x:x:x:d.d.d.d<br/>x - [0 to FFFF]H<br/>d - [0 to 255]D</p>                                                                                                                                                                                                                                                                                                                                                                                                               |

## statistics

|                |                                                                                                                                           |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>statistics</b> [ <b>sap</b> <i>sap-id</i>   <b>sdp</b> <i>sdp-id:vc-id</i>   <b>interface</b> <i>ip-address</i>   <i>ip-int-name</i> ] |
| <b>Context</b> | clear>service>id>dhcp                                                                                                                     |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command clears DHCP statistics.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <p><i>sap-id</i> — Clears the specified SAP information.</p> <p><i>sdp-id</i> — The specified SDP to be cleared.</p> <p><b>Values</b> 1 to 17407</p> <p><i>vc-id</i> — The virtual circuit ID on the SDP ID to be cleared.</p> <p><b>Values</b> 1 to 4294967295</p> <p><i>ip-address</i> — The interface IP address.</p> <p><b>Values</b> a.b.c.d</p> <p><i>ip-int-name</i> — The interface name. 32 characters maximum.</p> |

### 2.6.2.3 IES Debug Commands

#### host-connectivity-verify

|                    |                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] host-connectivity-verify</b>                                                                                                                          |
| <b>Context</b>     | debug>service>id                                                                                                                                              |
| <b>Description</b> | <p>This command enables Subscriber Host Connectivity Verification (SHCV) debugging.</p> <p>The <b>no</b> form of the command disables the SHCV debugging.</p> |

#### ip

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ip ip-address</b>                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | debug>service>id>host-connectivity-verify                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular IP address.                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <p><i>ip-address</i> — The IP address of the IP interface. The <i>ip-address</i> portion of the <b>address</b> command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> |

## mac

|                    |                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] mac</b> <i>ieee-address</i>                                                                                                                                                                                                                                                 |
| <b>Context</b>     | debug>service>id>host-connectivity-verify                                                                                                                                                                                                                                           |
| <b>Description</b> | This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular MAC address.                                                                                                                                                                         |
| <b>Parameters</b>  | <i>ieee-address</i> — 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. |

## sap

|                    |                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] sap</b> <i>sap-id</i>                                                                       |
| <b>Context</b>     | debug>service>id>host-connectivity-verify                                                           |
| <b>Description</b> | This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular SAP. |
| <b>Parameters</b>  | <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition.               |

## packets

|                    |                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] packets</b><br><b>[no] packets interface</b> <i>ip-int-name</i> [ <b>vrid</b> <i>virtual-router-id</i> ]<br><b>[no] packets interface</b> <i>ip-int-name</i> <b>vrid</b> <i>virtual-router-id</i> <b>ipv6</b> |
| <b>Context</b>     | debug>router>vrrp                                                                                                                                                                                                     |
| <b>Description</b> | This command enables or disables debugging for VRRP packets.                                                                                                                                                          |
| <b>Parameters</b>  | <i>ip-int-name</i> — Specifies the interface name. 32 characters maximum.<br><i>virtual-router-id</i> — Specifies the router ID.<br><b>Values</b> 1 to 255<br><b>ipv6</b> — Displays IPv6 information.                |

## events

|               |                                                                                                                |
|---------------|----------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>[no] events</b><br><b>[no] events interface</b> <i>ip-int-name</i> [ <b>vrid</b> <i>virtual-router-id</i> ] |
|---------------|----------------------------------------------------------------------------------------------------------------|

---

**[no] events interface** *ip-int-name* **vrid** *virtual-router-id* **ipv6**

- Context** debug>router>vrrp
- Description** This command enables or disables debugging for VRRP events.
- Parameters** *ip-int-name* — Specifies the interface name.  
*virtual-router-id* — Specifies the router ID.  
**Values** 1 to 255  
**ipv6** — Displays IPv6 information.

## 2.6.2.4 IES Monitor Commands

### instance

- Syntax** **instance interface** *interface-name* **vr-id** *virtual-router-id* [**ipv6**] [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]
- Context** monitor>router>vrrp
- Description** This command enables monitoring for statistics for VRRP instances.
- Parameters** *interface-name* — The name of the existing IP interface on which VRRP is configured. 32 characters maximum.  
*virtual-router-id* — The virtual router ID for the existing IP interface, expressed as a decimal integer.  
**Values** 1 to 255  
**ipv6** — Displays IPv6 information.  
*seconds* — Configures the interval for each display in seconds.  
**Values** 3 to 60  
**Default** 10  
*repeat* — Configures how many times the command is repeated.  
**Values** 1 to 999  
**Default** 10  
**absolute** — Raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.  
**rate** — Rate-per-second for each statistic is displayed instead of the delta.



---

## 3 Virtual Private Routed Network Service

### 3.1 VPRN Service Overview


RFC 2547b is an extension to the original RFC 2547, *BGP/MPLS VPNs*, which details a method of distributing routing information using BGP and MPLS 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 and handles the possibility of IP address overlap.

The service provider uses BGP to exchange the routes of a particular VPN among the PE routers that are attached to that VPN. This is done in a way which ensures that routes from different VPNs remain distinct and separate, even if two VPNs have an overlapping address space. The PE routers peer with locally connected CE routers and exchange routes with other PE routers in order to provide end-to-end connectivity between CEs belonging to a given VPN. Since the CE routers do not peer with each other there is no overlay visible to the CEs.

When BGP distributes a VPN route it also distributes an MPLS label for that route. On an SR series router, the method of allocating a label to a VPN route depends on the VPRN label mode and the configuration of the VRF export policy. SR series routers support three label allocation methods: label per VRF, label per next hop, and label per prefix.

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 one or additional MPLS labels or GRE tunnel header so that it gets tunneled across the backbone to the proper PE router. Each route exchanged by the MP-BGP protocol includes a route distinguisher (RD), which identifies the VPRN association. Thus the backbone core routers do not need to know the VPN routes. [Figure 9](#) displays a VPRN network diagram example.

**Figure 9 Virtual Private Routed Network**

### 3.1.1 Routing Prerequisites

RFC4364 requires the following features:

- multi-protocol extensions to BGP
- extended BGP community support
- BGP capability negotiation

Tunneling protocol options are as follows:

- Label Distribution Protocol (LDP)
- MPLS RSVP-TE tunnels
- Generic Router Encapsulation (GRE) tunnels
- BGP route tunnel (RFC3107)

### 3.1.2 Core MP-BGP Support


BGP is used with BGP extensions mentioned in [Routing Prerequisites](#) to distribute VPRN routing information across the service provider's network.

BGP was initially designed to distribute IPv4 routing information. Therefore, multi-protocol extensions and the use of a VPN-IP 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. A VPN-IPv6 address is a 24-byte value consisting of the 8-byte RD and 16-byte IPv6 address prefix. Service providers typically assign one or a small number of RDs per VPN service network-wide.

### 3.1.3 Route Distinguishers

The route distinguisher (RD) is an 8-byte value consisting of two major fields, the **Type** field and **Value** field. The **Type** field determines how the **Value** field should be interpreted. The 7750 SR and 7950 XRS implementation supports the three (3) **Type** values as defined in the standard.

**Figure 10      Route Distinguisher**



The three Type values are:

- Type 0: Value Field — Administrator subfield (2 bytes)  
Assigned number subfield (4 bytes)

The administrator field must contain an AS number (using private AS numbers is discouraged). The Assigned field contains a number assigned by the service provider.

- Type 1: Value Field — Administrator subfield (4 bytes)  
Assigned number subfield (2 bytes)

The administrator field must contain an IP address (using private IP address space is discouraged). The Assigned field contains a number assigned by the service provider.

- Type 2: Value Field — Administrator subfield (4 bytes)  
Assigned number subfield (2 bytes)


The administrator field must contain a 4-byte AS number (using private AS numbers is discouraged). The Assigned field contains a number assigned by the service provider.

### 3.1.3.1 eiBGP Load Balancing

eiBGP load balancing allows a route to have multiple nexthops of different types, using both IPv4 nexthops and MPLS LSPs simultaneously.

Figure 11 displays a basic topology that could use eiBGP load balancing. In this topology CE1 is dual homed and thus reachable by two separate PE routers. CE 2 (a site in the same VPRN) is also attached to PE1. With eiBGP load balancing, PE1 will utilize its own local IPv4 nexthop as well as the route advertised by MP-BGP, by PE2.


**Figure 11 Basic eiBGP Topology**



Another example displayed in Figure 12 shows an extra net VPRN (VRF). The traffic ingressing the PE that should be load balanced is part of a second VPRN and the route over which the load balancing is to occur is part of a separate VPRN instance and are leaked into the second VPRN by route policies.

Here, both routes can have a source protocol of VPN-IPv4 but one will still have an IPv4 nexthop and the other can have a VPN-IPv4 nexthop pointing out a network interface. Traffic will still be load balanced (if eiBGP is enabled) as if only a single VRF was involved.

**Figure 12 Extranet Load Balancing**



Traffic will be load balanced across both the IPv4 and VPN-IPv4 next hops. This helps to use all available bandwidth to reach a dual-homed VPRN.

### 3.1.4 Route Reflector

The use of Route Reflectors is supported in the service provider core. Multiple sets of route reflectors can be used for different types of BGP routes, including IPv4 and VPN-IPv4 as well as multicast and IPv6 (multicast and IPv6 apply to the 7750 SR only).

### 3.1.5 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
- E-BGP
- RIP
- OSPF
- OSPF3

Each protocol provides controls to limit the number of routes learned from each CE router.

---

### 3.1.5.1 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 certain routes before they are injected into the local VRF.

Route redistribution from the local VRF to CE-to-PE routing protocols is to be controlled via the route policies in each routing protocol instance, in the same manner that is used by the base router instance.


The advertisement or redistribution of routing information from the local VRF to or from the MP-BGP instance is specified per VRF and is controlled by VRF route target associations or by VRF route policies.

VPN-IP routes imported into a VPRN, have the protocol **type bgp-vpn** to denote that it is an VPRN route. This can be used within the route policy match criteria.

### 3.1.5.2 CPE Connectivity Check


Static routes are used within many IES services and VPRN services. Unlike dynamic routing protocols, there is no way to change the state of routes based on availability information for the associated CPE. CPE connectivity check adds flexibility so that unavailable destinations will be removed from the VPRN routing tables dynamically and minimize wasted bandwidth. [Figure 13](#) shows a setup with a directly connected IP target and [Figure 14](#) shows a setup with multiple hops to an IP target.

**Figure 13 Directly Connected IP Target**



Fig\_18

**Figure 14 Multiple Hops to IP Target**



Fig\_19

The availability of the far-end static route is monitored through periodic polling. The polling period is configured. If the poll fails a specified number of sequential polls, the static route is marked as inactive.

Either ICMP ping or unicast ARP mechanism can be used to test the connectivity. ICMP ping is preferred.

If the connectivity check fails and the static route is deactivated, the router will continue to send polls and re-activate any routes that are restored.

---

## 3.1.6 Constrained Route Distribution (RT Constraint)

### 3.1.6.1 Constrained VPN Route Distribution Based on Route Targets

Constrained Route Distribution (or RT Constraint) is a mechanism that allows a router to advertise Route Target membership information to its BGP peers to indicate interest in receiving only VPN routes tagged with specific Route Target extended communities. Upon receiving this information, peers restrict the advertised VPN routes to only those requested, minimizing control plane load in terms of protocol traffic and possibly also RIB memory.

The Route Target membership information is carried using MP-BGP, using an AFI value of 1 and SAFI value of 132. In order for two routers to exchange RT membership NLRI they must advertise the corresponding AFI/SAFI to each other during capability negotiation. The use of MP-BGP means RT membership NLRI are propagated, loop-free, within an AS and between ASes using well-known BGP route selection and advertisement rules.

ORF can also be used for RT-based route filtering, but ORF messages have a limited scope of distribution (to direct peers) and therefore do not automatically create pruned inter-cluster and inter-AS route distribution trees.

### 3.1.6.2 Configuring the Route Target Address Family

RT Constraint is supported only by the base router BGP instance. When the **family** command at the BGP router group or neighbor CLI context includes the **route-target** keyword, the RT Constraint capability is negotiated with the associated set of EBGP and IBGP peers.

ORF is mutually exclusive with RT Constraint on a particular BGP session. The CLI will not attempt to block this configuration, but if both capabilities are enabled on a session, the ORF capability will not be included in the OPEN message sent to the peer.



### 3.1.6.3 Originating RT Constraint Routes

When the base router has one or more RTC peers (BGP peers with which the RT Constraint capability has been successfully negotiated), one RTC route is created for each RT extended community imported into a locally-configured L2 VPN or L3 VPN service. These imported route targets are configured in the following contexts:

- `config>service>vpn`
- `config>service>vpn>mvpn`

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, a 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 **config>router>autonomous-system** command) and 0 or 16-64 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. This is done using the new **default-route-target** group/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.

### 3.1.6.4 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 applies:

- The prefix length is 1-31.
- The prefix length is 33-47.
- The prefix length is 48-96 and the 16 most-significant bits are not 0x0002, 0x0102 or 0x0202.

If multiple RTC routes are received for the same prefix value then 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 or an ASBR does not re-advertise the best RTC route to any RTC peer due to standard IBGP split horizon rules.
- A route reflector that receives its best RTC route for a prefix from a client peer 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 (i) sets the ORIGINATOR\_ID to its own router ID and (ii) 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. No aggregation of RTC routes is 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. No aggregation of RTC routes is supported.



**Note:** These advertisement rules do not handle hierarchical RR topologies properly. This is a limitation of the current RT constraint standard.

### 3.1.6.5 Using RT Constraint Routes

In general (ignoring IBGP-to-IBGP rules, Add-Path, Best-external, etc.), 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 some prefix/NLRI from being advertised to specific peers. These export policies may be configured statically or created dynamically based on use of ORF or RT constraint with a peer. ORF and RT Constraint are mutually exclusive on a session.

When RT Constraint is configured on a session that also supports VPN address families using route targets (that is: vpn-ipv4, vpn-ipv6, l2-vpn, mvpn-ipv4, mvpn-ipv6, mcast-vpn-ipv4 or evpn), the advertisement of the VPN routes is affected as follows:

- When the session comes up, the advertisement of the VPN routes is delayed for a short while to allow RTC routes to be received from the peer.
- After the initial delay, the received RTC routes are analyzed and acted upon. If S1 is the set of routes previously advertised to the peer and S2 is the set of routes that should be advertised based on the most recent received RTC routes then:
  - Set of routes in S1 but not in S2 should be withdrawn immediately (subject to MRAI).
  - 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 is the set that:
  - a. are eligible for advertisement to P1 per BGP route advertisement rules AND
  - b. have not been rejected by manually configured export policies AND
  - c. have not been advertised to the peer



**Note:** This applies whether or not P1 advertised the best route for the default RTC prefix.

In this context, a default RTC route is any of the following:

1. a route with NLRI length = zero
2. a route with NLRI value = origin AS and NLRI length = 32
3. 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:
    - a. are eligible for advertisement to I1 per BGP route advertisement rules AND
    - b. have not been rejected by manually configured export policies AND
    - c. carry at least one route target extended community with value A2 in the n most significant bits AND
    - d. have not been advertised to the peer



**Note:** This applies whether or not I1 advertised the best route for A.

- If the best RTC route for a prefix A (origin-AS = A1, RT = A2/n,  $n > 48$ ) is received from an IBGP peer I1 in autonomous system B, the VPN routes that are advertised to I1 is the set that:
  - a. are eligible for advertisement to I1 per BGP route advertisement rules AND
  - b. have not been rejected by manually configured export policies AND
  - c. carry at least one route target extended community with value A2 in the n most significant bits AND
  - d. 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:
  - a. are eligible for advertisement to E1 per BGP route advertisement rules AND
  - b. have not been rejected by manually configured export policies AN
  - c. carry at least one route target extended community with value A2 in the n most significant bits AND
  - d. have not been advertised to the peer



**Note:** This applies only if E1 advertised the best route for A.

### 3.1.7 BGP Fast Reroute in a VPRN

BGP fast reroute is a feature that brings together indirection techniques in the forwarding plane and pre-computation of BGP backup paths in the control plane to support fast reroute of BGP traffic around unreachable/failed next-hops. In a VPRN context BGP fast reroute is supported using unlabeled IPv4, unlabeled IPv6, VPN-IPv4, and VPN-IPv6 VPN routes. The supported VPRN scenarios are outlined in [Table 31](#).

BGP fast reroute information specific to the base router BGP context is described in the BGP Fast Reroute section of the *7450 ESS, 7750 SR, 7950 XRS, and VSR Unicast Routing Protocols Guide*.

**Table 31 BGP Fast Reroute Scenarios (VPRN Context)**

| Ingress Packet    | Primary Route                                                             | Backup Route                                                               | Prefix Independent Convergence |
|-------------------|---------------------------------------------------------------------------|----------------------------------------------------------------------------|--------------------------------|
| IPv4 (ingress PE) | IPv4 route with next-hop A resolved by an IPv4 route                      | IPv4 route with next-hop B resolved by an IPv4 route                       | Yes                            |
| IPv4 (ingress PE) | VPN-IPv4 route with next-hop A resolved by a GRE, LDP, RSVP or BGP tunnel | VPN-IPv4 route with next-hop A resolved by a GRE, LDP, RSVP or BGP tunnel  | Yes                            |
| MPLS (egress PE)  | IPv4 route with next-hop A resolved by an IPv4 route                      | IPv4 route with next-hop B resolved by an IPv4 route                       | Yes                            |
| MPLS (egress PE)  | IPv4 route with next-hop A resolved by an IPv4 route                      | VPN-IPv4 route* with next-hop B resolved by a GRE, LDP, RSVP or BGP tunnel | Yes                            |
| IPv6 (ingress PE) | IPv6 route with next-hop A resolved by an IPv6 route                      | IPv6 route with next-hop B resolved by an IPv6 route                       | Yes                            |
| IPv6 (ingress PE) | VPN-IPv6 route with next-hop A resolved by a GRE, LDP, RSVP or BGP tunnel | VPN-IPv6 route with next-hop B resolved by a GRE, LDP, RSVP or BGP tunnel  | Yes                            |
| MPLS (egress)     | IPv6 route with next-hop A resolved by an IPv6 route                      | IPv6 route with next-hop B resolved by an IPv6 route                       | Yes                            |

**Table 31 BGP Fast Reroute Scenarios (VPRN Context) (Continued)**

| Ingress Packet | Primary Route                                        | Backup Route | Prefix Independent Convergence                                                                                                                                                                                |
|----------------|------------------------------------------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MPLS (egress)  | IPv6 route with next-hop A resolved by an IPv6 route | Yes          | VPRN label mode must be VRF. VPRN must export its VPN-IP routes with RD $\neq$ y. For the best performance the backup next-hop must advertise the same VPRN label value with all routes (e.g. per VRF label). |

### 3.1.7.1 BGP Fast Reroute in a VPRN Configuration

In a VPRN context, BGP fast reroute is optional and must be enabled. Fast reroute can be applied to all IPv4 prefixes, all IPv6 prefixes, all IPv4 and IPv6 prefixes, or to a specific set of IPv4 and IPv6 prefixes.

If all IP prefixes require backup path protection, use a combination of the BGP instance-level **backup-path** and VPRN-level **enable-bgp-vpn-backup** commands. The VPRN BGP **backup-path** command enables BGP fast reroute for all IPv4 prefixes and/or all IPv6 prefixes that have a best path through a VPRN BGP peer. The VPRN-level **enable-bgp-vpn-backup** command enables BGP fast reroute for all IPv4 prefixes and/or all IPv6 prefixes that have a best path through a remote PE peer.

If only some IP prefixes require backup path protection, use route policies to apply the **install-backup-path** action to the best paths of the IP prefixes requiring protection. See the “BGP Fast Reroute” section of the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Unicast Routing Protocols Guide* for more information.

### 3.1.8 BGP Best-External in a VPRN Context

If two or more PE routers connect to a multi-homed site and learn routes for a common set of IP prefixes from that site, then the failure of one of the PE routers or a PE-CE link can be handled by rerouting the traffic over the alternate paths. The traffic failover time in this situation can be reduced if all the PE routers have advance knowledge of the potential backup paths and do not have to wait for BGP route advertisements and/or withdrawals to reprogram their forwarding tables. This can be challenging with normal BGP procedures because a PE router is not allowed to

advertise, to other PE routers, a BGP route that it has learned from a connected CE device if that route is not its active route for the destination in the route table. If the multi-homing scenario calls for all traffic destined for an IP prefix to be carried over a preferred primary path (passing through PE1-CE1 for example), then all other PE routers (PE2, PE3, and so on) will have that VPN route as their active route for the destination, and they will not be able to advertise their own routes for the same IP prefix.

The SR OS supports a VPRN feature, configured using the **export-inactive-bgp** command, that resolves the issue described above. When a VPRN is configured with this command, it is allowed to advertise (as a VPN-IP route towards other PEs) its best CE-BGP route for an IP prefix, even when that CE-BGP route is inactive in the route table due to the presence of a more-preferred VPN-IP route from another PE. In order for the CE-BGP route to be advertised, the CE-BGP route must be accepted by the VRF export policy. When a VPN-IP route is advertised due to the **export-inactive-bgp** command, the label carried in the route is a per-next-hop label corresponding to the next-hop IP address of the CE-BGP route, or a per-prefix label; this helps avoid packet looping issues due to unsynchronized IP FIBs.

When a PE router that advertised a backup path for an IP prefix receives a withdrawal for the VPN-IP route that it was using as the primary/active route, its backup path may be promoted to the primary path; that is, the CE-BGP route may become the active route for the destination. In this case, the PE router is required to re-advertise the VPN-IP route with a per-VRF label if that is the default allocation policy and there is no label-per-prefix policy override. It will take some time for the new VPN-IP route to reach all the ingress routers and for them to update their forwarding tables. In the meantime, traffic will continue to be received with the old per-next-hop label. The egress PE will drop this in-flight traffic unless **label retention** is configured using the **bgp-labels-hold-timer** command in the **config>router>mpls-labels** context. This command configures a delay (in seconds) between the withdrawal of a VPN-IP route with a per-next-hop label and the deletion of the corresponding label forwarding entry in the IOM. The value of **bgp-labels-hold-timer** should be large enough to account for the propagation delay of the route withdrawal to all the ingress routers.

---

## 3.2 VPRN Features

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

### 3.2.1 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:

- VRRP
- Cflowd
- Secondary IP addresses
- ICMP Options

Configuration options found on core IP interfaces not supported on VPRN IP interfaces are:

- NTP broadcast receipt

#### 3.2.1.1 QoS Policy Propagation Using BGP (QPPB)

This section discusses QPPB as it applies to VPRN, IES, and router interfaces. Refer to the QoS Policy Propagation Using BGP (QPPB) section and the IP Router Configuration section in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide*.

The QoS Policy Propagation Using BGP (QPPB) feature applies only to the 7450 ESS and 7750 SR.

QoS policy propagation using BGP (QPPB) is a feature that allows a route to be installed in the routing table with a forwarding-class and priority so that packets matching the route can receive the associated QoS. The forwarding-class and priority associated with a BGP route are set using BGP import route policies. In the industry, this feature is called QPPB, and even though the feature name refers to BGP specifically. On SR OS, QPPB is supported for BGP (IPv4, IPv6, VPN-IPv4, VPN-IPv6), RIP and static routes.



While SAP ingress and network QoS policies can achieve the same end result as QPPB, the effort involved in creating the QoS policies, keeping them up-to-date, and applying them across many nodes is much greater than with QPPB. This is due to assigning a packet, arriving on a particular IP interface, to a specific forwarding-class and priority/profile, based on the source IP address or destination IP address of the packet. In a typical application of QPPB, a BGP route is advertised with a BGP community attribute that conveys a particular QoS. Routers that receive the advertisement accept the route into their routing table and set the forwarding-class and priority of the route from the community attribute.

### **3.2.1.2 QPPB Applications**

There are two typical applications of QPPB:

1. coordination of QoS policies between different administrative domains, and
2. traffic differentiation within a single domain, based on route characteristics.

### **3.2.1.3 Inter-AS Coordination of QoS Policies**

The operator of an administrative domain A can use QPPB to signal to a peer administrative domain B that traffic sent to certain prefixes advertised by domain A should receive a particular QoS treatment in domain B. More specifically, an ASBR of domain A can advertise a prefix XYZ to domain B and include a BGP community attribute with the route. The community value implies a particular QoS treatment, as agreed by the two domains (in their peering agreement or service level agreement, for example). When the ASBR and other routers in domain B accept and install the route for XYZ into their routing table, they apply a QoS policy on selected interfaces that classifies traffic towards network XYZ into the QoS class implied by the BGP community value.

QPPB may also be used to request that traffic sourced from certain networks receive appropriate QoS handling in downstream nodes that may span different administrative domains. This can be achieved by advertising the source prefix with a BGP community, as discussed above. However, in this case other approaches are equally valid, such as marking the DSCP or other CoS fields based on source IP address so that downstream domains can take action based on a common understanding of the QoS treatment implied by different DSCP values.


In the above examples, coordination of QoS policies using QPPB could be between a business customer and its IP VPN service provider, or between one service provider and another.

### 3.2.1.4 Traffic Differentiation Based on Route Characteristics

There may be times when a network operator wants to provide differentiated service to certain traffic flows within its network, and these traffic flows can be identified with known routes. For example, the operator of an ISP network may want to give priority to traffic originating in a particular ASN (the ASN of a content provider offering over-the-top services to the ISP's customers), following a certain AS\_PATH, or destined for a particular next-hop (remaining on-net vs. off-net).

Figure 15 shows an example of an ISP that has an agreement with the content provider managing AS300 to provide traffic sourced and terminating within AS300 with differentiated service appropriate to the content being transported. In this example we presume that ASBR1 and ASBR2 mark the DSCP of packets terminating and sourced, respectively, in AS300 so that other nodes within the ISP's network do not need to rely on QPPB to determine the correct forwarding-class to use for the traffic. The DSCP or other CoS markings could be left unchanged in the ISP's network and QPPB used on every node.

**Figure 15 Use of QPPB to Differentiate Traffic in an ISP Network**



### 3.2.1.5 QPPB

There are two main aspects of the QPPB feature on the 7450 ESS and 7750 SR:

- the ability to associate a forwarding-class and priority with certain routes in the routing table, and
- the ability to classify an IP packet arriving on a particular IP interface to the forwarding-class and priority associated with the route that best matches the packet.

### 3.2.1.6 Associating an FC and Priority with a Route

This feature uses a command in the route-policy hierarchy to set the forwarding class and optionally the priority associated with routes accepted by a route-policy entry. The command has the following structure:

**fc** *fc-name* [**priority** {**low** | **high**}]

The use of this command is illustrated by the following example:

```
config>router>policy-options
begin
community gold members 300:100
policy-statement qppb_policy
entry 10
from
protocol bgp
community gold
exit
action accept
fc hl priority high
exit
exit
exit
commit
```

The **fc** command is supported with all existing from and to match conditions in a route policy entry and with any action other than reject, it is supported with next-entry, next-policy and accept actions. If a next-entry or next-policy action results in multiple matching entries then the last entry with a QPPB action determines the forwarding class and priority.

A route policy that includes the **fc** command in one or more entries can be used in any import or export policy but the **fc** command has no effect except in the following types of policies:

- VRF import policies:
  - config>service>vprn>vrf-import
- BGP import policies:
  - config>router>bgp>import
  - config>router>bgp>group>import
  - config>router>bgp>group>neighbor>import
  - config>service>vprn>bgp>import
  - config>service>vprn>bgp>group>import
  - config>service>vprn>bgp>group>neighbor>import
- RIP import policies:

- config>router>rip>import
- config>router>rip>group>import
- config>router>rip>group>neighbor>import
- config>service>vprn>rip>import
- config>service>vprn>rip>group>import
- config>service>vprn>rip>group>neighbor>import

As evident from above, QPPB route policies support routes learned from RIP and BGP neighbors of a VPRN as well as for routes learned from RIP and BGP neighbors of the base/global routing instance.

QPPB is supported for BGP routes belonging to any of the address families listed below:

- IPv4 (AFI=1, SAFI=1)
- IPv6 (AFI=2, SAFI=1)
- VPN-IPv4 (AFI=1, SAFI=128)
- VPN-IPv6 (AFI=2, SAFI=128)

A VPN-IP route may match both a VRF import policy entry and a BGP import policy entry (if vpn-apply-import is configured in the base router BGP instance). In this case the VRF import policy is applied first and then the BGP import policy, so the QPPB QoS is based on the BGP import policy entry.

This feature also introduces the ability to associate a forwarding-class and optionally priority with IPv4 and IPv6 static routes. This is achieved by specifying the forwarding-class within the static-route-entry next-hop or indirect context.

Priority is optional when specifying the forwarding class of a static route, but once configured it can only be deleted and returned to unspecified by deleting the entire static route.

### 3.2.1.7 Displaying QoS Information Associated with Routes

The following commands are enhanced to show the forwarding-class and priority associated with the displayed routes:

- show router route-table
- show router fib
- show router bgp routes
- show router rip database

- show router static-route

This feature uses a **qos** keyword to the **show>router>route-table** command. When this option is specified the output includes an additional line per route entry that displays the forwarding class and priority of the route. If a route has no fc and priority information then the third line is blank. The following CLI shows an example:

**show router route-table [family] [ip-prefix[/prefix-length]] [longer | exact]  
[protocol protocol-name] qos**

An example output of this command is shown below:

```
A:Dut-A# show router route-table 10.1.5.0/24 qos
=====
Route Table (Router: Base)
=====
Dest Prefix Type Proto Age Pref
 Next Hop[Interface Name] Metric
 QoS

10.1.5.0/24 Remote BGP 15h32m52s 0
 PE1_to_PE2
 h1, high

No. of Routes: 1
=====
A:Dut-A#
```

### 3.2.1.8 Enabling QPPB on an IP interface

To enable QoS classification of ingress IP packets on an interface based on the QoS information associated with the routes that best match the packets the **qos-route-lookup** command is necessary in the configuration of the IP interface. The **qos-route-lookup** command has parameters to indicate whether the QoS result is based on lookup of the source or destination IP address in every packet. There are separate qos-route-lookup commands for the IPv4 and IPv6 packets on an interface, which allows QPPB to be enabled for IPv4 only, IPv6 only, or both IPv4 and IPv6. Current QPPB based on a source IP address is not supported for IPv6 packets nor is it supported for ingress subscriber management traffic on a group interface.

The qos-route-lookup command is supported on the following types of IP interfaces:

- base router network interfaces (config>router>interface)
- VPRN SAP and spoke SDP interfaces (config>service>vprn>interface)
- VPRN group-interfaces (config>service>vprn>sub-if>grp-if)
- IES SAP and spoke SDP interfaces (config>service>ies>interface)
- IES group-interfaces (config>service>ies>sub-if>grp-if)

When the `qos-route-lookup` command with the destination parameter is applied to an IP interface and the destination address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the SAP-Ingress or network qos policy associated with the IP interface. If the destination address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the SAP-Ingress or network qos policy.

Similarly, when the `qos-route-lookup` command with the source parameter is applied to an IP interface and the source address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the SAP-Ingress or network qos policy associated with the IP interface. If the source address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the SAP-Ingress or network qos policy.

Currently, QPPB is not supported for ingress MPLS traffic on network interfaces or on CsC PE'-CE' interfaces (`config>service>vpn>nw-if`).

### 3.2.1.9 QPPB When Next-Hops are Resolved by QPPB Routes

In some circumstances (IP VPN inter-AS model C, Carrier Supporting Carrier, indirect static routes, etc.) an IPv4 or IPv6 packet may arrive on a QPPB-enabled interface and match a route A1 whose next-hop N1 is resolved by a route A2 with next-hop N2 and perhaps N2 is resolved by a route A3 with next-hop N3, etc. The QPPB result is based only on the forwarding-class and priority of route A1. If A1 does not have a forwarding-class and priority association then the QoS classification is not based on QPPB, even if routes A2, A3, etc. have forwarding-class and priority associations.

### 3.2.1.10 QPPB and Multiple Paths to a Destination

When ECMP is enabled some routes may have multiple equal-cost next-hops in the forwarding table. When an IP packet matches such a route the next-hop selection is typically based on a hash algorithm that tries to load balance traffic across all the next-hops while keeping all packets of a given flow on the same path. The QPPB configuration model described in [Associating an FC and Priority with a Route](#) allows different QoS information to be associated with the different ECMP next-hops of a route. The forwarding-class and priority of a packet matching an ECMP route is based on the particular next-hop used to forward the packet.

When BGP FRR is enabled some BGP routes may have a backup next-hop in the forwarding table in addition to the one or more primary next-hops representing the equal-cost best paths allowed by the ECMP/multipath configuration. When an IP packet matches such a route a reachable primary next-hop is selected (based on the hash result) but if all the primary next-hops are unreachable then the backup next-hop is used. The QPPB configuration model described in [Associating an FC and Priority with a Route](#) allows the forwarding-class and priority associated with the backup path to be different from the QoS characteristics of the equal-cost best paths. The forwarding class and priority of a packet forwarded on the backup path is based on the **fc** and priority of the backup route.

### 3.2.1.11 QPPB and Policy-Based Routing

When an IPv4 or IPv6 packet with destination address X arrives on an interface with both QPPB and policy-based-routing enabled:

- There is no QPPB classification if the IP filter action redirects the packet to a directly connected interface, even if X is matched by a route with a forwarding-class and priority
- QPPB classification is based on the forwarding-class and priority of the route matching IP address Y if the IP filter action redirects the packet to the indirect next-hop IP address Y, even if X is matched by a route with a forwarding-class and priority.

### 3.2.1.12 QPPB and GRT Lookup

Source-address based QPPB is not supported on any SAP or spoke SDP interface of a VPRN configured with the **grt-lookup** command.

### 3.2.1.13 QPPB Interaction with SAP Ingress QoS Policy

When QPPB is enabled on a SAP IP interface the forwarding class of a packet may change from **fc1**, the original **fc** determined by the SAP ingress QoS policy to **fc2**, the new **fc** determined by QPPB. In the ingress datapath SAP ingress QoS policies are applied in the first P chip and route lookup/QPPB occurs in the second P chip. This has the implications listed below:

- Ingress remarking (based on profile state) is always based on the original **fc** (**fc1**) and sub-class (if defined).

- The profile state of a SAP ingress packet that matches a QPPB route depends on the configuration of **fc2** only. If the de-1-out-profile flag is enabled in **fc2** and **fc2** is not mapped to a priority mode queue then the packet will be marked out of profile if its DE bit = 1. If the profile state of **fc2** is explicitly configured (in or out) and **fc2** is not mapped to a priority mode queue then the packet is assigned this profile state. In both cases there is no consideration of whether or not **fc1** was mapped to a priority mode queue.
- The priority of a SAP ingress packet that matches a QPPB route depends on several factors. If the de-1-out-profile flag is enabled in **fc2** and the DE bit is set in the packet then priority will be low regardless of the QPPB priority or **fc2** mapping to profile mode queue, priority mode queue or policer. If **fc2** is associated with a profile mode queue then the packet priority will be based on the explicitly configured profile state of **fc2** (in profile = high, out profile = low, undefined = high), regardless of the QPPB priority or **fc1** configuration. If **fc2** is associated with a priority mode queue or policer then the packet priority will be based on QPPB (unless DE=1), but if no priority information is associated with the route then the packet priority will be based on the configuration of **fc1** (if **fc1** mapped to a priority mode queue then it is based on DSCP/IP prec/802.1p and if **fc1** mapped to a profile mode queue then it is based on the profile state of **fc1**).

Table 32 summarizes these interactions.

**Table 32 QPPB Interactions with SAP Ingress QoS**

| Original FC object mapping | New FC object mapping | Profile                                    | Priority (drop preference)                                                                                                              | DE=1 override    | In/out of profile marking      |
|----------------------------|-----------------------|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|------------------|--------------------------------|
| Profile mode queue         | Profile mode queue    | From new base FC unless overridden by DE=1 | From QPPB, unless packet is marked in or out of profile in which case follows profile. Default is high priority.                        | From new base FC | From original FC and sub-class |
| Priority mode queue        | Priority mode queue   | Ignored                                    | If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then from original dot1p/exp/DSCP mapping or policy default. | From new base FC | From original FC and sub-class |



**Table 32 QPPB Interactions with SAP Ingress QoS (Continued)**

| Original FC object mapping | New FC object mapping | Profile                                    | Priority (drop preference)                                                                                                              | DE=1 override    | In/out of profile marking      |
|----------------------------|-----------------------|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|------------------|--------------------------------|
| Policer                    | Policer               | From new base FC unless overridden by DE=1 | If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then from original dot1p/exp/DSCP mapping or policy default. | From new base FC | From original FC and sub-class |
| Priority mode queue        | Policer               | From new base FC unless overridden by DE=1 | If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then from original dot1p/exp/DSCP mapping or policy default. | From new base FC | From original FC and sub-class |
| Policer                    | Priority mode queue   | Ignored                                    | If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then from original dot1p/exp/DSCP mapping or policy default. | From new base FC | From original FC and sub-class |
| Profile mode queue         | Priority mode queue   | Ignored                                    | If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then follows original FC's profile mode rules.               | From new base FC | From original FC and sub-class |
| Priority mode queue        | Profile mode queue    | From new base FC unless overridden by DE=1 | From QPPB, unless packet is marked in or out of profile in which case follows profile. Default is high priority.                        | From new base FC | From original FC and sub-class |
| Profile mode queue         | Policer               | From new base FC unless overridden by DE=1 | If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then follows original FC's profile mode rules.               | From new base FC | From original FC and sub-class |

**Table 32 QPPB Interactions with SAP Ingress QoS (Continued)**

| Original FC object mapping | New FC object mapping | Profile                                    | Priority (drop preference)                                                                                       | DE=1 override    | In/out of profile marking      |
|----------------------------|-----------------------|--------------------------------------------|------------------------------------------------------------------------------------------------------------------|------------------|--------------------------------|
| Policer                    | Profile mode queue    | From new base FC unless overridden by DE=1 | From QPPB, unless packet is marked in or out of profile in which case follows profile. Default is high priority. | From new base FC | From original FC and sub-class |

### 3.2.1.14 Object Grouping and State Monitoring

This feature introduces a generic operational group object which associates different service endpoints (pseudowires and SAPs) located in the same or in different service instances. The operational group status is derived from the status of the individual components using certain rules specific to the application using the concept. A number of other service entities, the monitoring objects, can be configured to monitor the operational group status and to perform certain actions as a result of status transitions. For example, if the operational group goes down, the monitoring objects will be brought down.

### 3.2.1.15 VPRN IP Interface Applicability

This concept is used by an IPv4 VPRN interface to affect the operational state of the IP interface monitoring the operational group. Individual SAP and spoke SDPs are supported as monitoring objects.

The following rules apply:

- An object can only belong to one group at a time.
- An object that is part of a group cannot monitor the status of a group.
- An object that monitors the status of a group cannot be part of a group.
- An operational group may contain any combination of member types: SAP or Spoke-SDPs.
- An operational group may contain members from different VPLS service instances.
- Objects from different services may monitor the oper-group.

There are two steps involved in enabling the functionality:

- Step 1.** Identify a set of objects whose forwarding state should be considered as a whole group then group them under an operational group using the **oper-group** command.
- Step 2.** Associate the IP interface to the oper-group using the **monitor-group** command

The status of the operational group (oper-group) is dictated by the status of one or more members according to the following rule:

- The oper-group goes down if all the objects in the oper-group go down. The oper-group comes up if at least one of the components is up.
- An object in the group is considered down if it is not forwarding traffic in at least one direction. That could be because the operational state is down or the direction is blocked through some validation mechanism.
- If a group is configured but no members are specified yet then its status is considered up.
- As soon as the first object is configured the status of the operational group is dictated by the status of the provisioned member(s).

The simple configuration below shows the oper-group g1, the VPLS SAP that is mapped to it and the IP interfaces in VPRN service 2001 monitoring the oper-group g1. This is example uses an R-VPLS context. The VPLS instance includes the **allow-ip-int-bind** and the **name v1**. The VPRN interface links to the VPLS using the **vpls v1** option. All commands are under the configuration service hierarchy.

To further explain the configuration. Oper-group g1 has a single SAP (1/1/1:2001) mapped to it and the IP interfaces in the VPRN service 2001 will derive its state from the state of oper-group g1.

```
oper-group g1 create

vpls 1 name "v1" customer 1 create
 allow-ip-int-bind
 stp
 shutdown
 exit
 sap 1/1/1:2001 create
 oper-group g1
 eth-cfm
 mep domain 1 association 1 direction down
ccm-enable
 no shutdown
 exit
 exit
 sap 1/1/2:2001 create
 exit
 sap 1/1/3:2001 create
 exit
no shutdown
```

```
vrpn 2001 customer 1 create
 interface "i2001" create
 address 10.1.1.1/24
 monitor-oper-group "g1"
 vpls "v1"
 exit
no shutdown
exit
```

## 3.2.2 Subscriber Interfaces

Subscriber interfaces are composed of a combination of two key technologies, subscriber interfaces and group interfaces. While the subscriber interface defines the subscriber subnets, the group interfaces are responsible for aggregating the SAPs.

Subscriber Interfaces apply only to the 7450 ESS and 7750 SR.

- Subscriber interface — an interface that allows the sharing of a subnet among one or many group interfaces in the routed CO model
- Group interface — aggregates multiple SAPs on the same port
- Redundant interfaces — a special spoke-terminated Layer 3 interface. It is used in a Layer 3 routed CO dual-homing configuration to shunt downstream (network to subscriber) to the active node for a given subscriber

## 3.2.3 SAPs

### 3.2.3.1 Encapsulations

The following SAP encapsulations are supported on the 7750 SR and 7950 XRS VPRN service:

- Ethernet null
- Ethernet dot1q
- SONET/SDH IPCP
- SONET/SDH ATM
- ATM - LLC SNAP or VC-MUX
- Cisco HDLC
- QinQ

- LAG
- Tunnel (IPSec or GRE)
- Frame Relay

### 3.2.3.2 ATM SAP Encapsulations for VPRN Services

The router supports ATM PVC service encapsulation for VPRN SAPs on the 7750 SR only. Both UNI and NNI cell formats are supported. The format is configurable on a SONET/SDH path basis. A path maps to an ATM VC. All VCs on a path must use the same cell format.

The following ATM encapsulation and transport modes are supported:

- RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*:
  - AAL5 LLC/SNAP IPv4 routed
  - AAL5 VC mux IPv4 routed
  - AAL5 LLC/SNAP IPv4 bridged
  - AAL5 VC mux IPv4 bridged

### 3.2.3.3 Pseudowire SAPs

Pseudowire SAPs are supported on VPRN interfaces for the 7750 SR in the same way as on IES interfaces.

## 3.2.4 QoS Policies

When applied to a VPRN SAP, service ingress QoS policies only create the unicast queues defined in the policy if PIM is not configured on the associated IP interface; if PIM is configured, the multipoint queues are applied as well.

With VPRN services, service egress QoS policies function as with other services where the class-based queues are created as defined in the policy.

Both Layer 2 or Layer 3 criteria can be used in the QoS policies for traffic classification in an VPRN.

## 3.2.5 Filter Policies

Ingress and egress IPv4 and IPv6 filter policies can be applied to VPRN SAPs.

## 3.2.6 DSCP Marking

Specific DSCP, forwarding class, and Dot1P parameters can be specified to be used by every protocol packet generated by the VPRN. This enables prioritization or de-prioritization of every protocol (as required). The markings effect a change in behavior on ingress when queuing. For example, if OSPF is not enabled, then traffic can be de-prioritized to best effort (be) DSCP. This change de-prioritizes OSPF traffic to the CPU complex.

DSCP marking for internally generated control and management traffic by marking the DSCP value should be used for the given application. This can be configured per routing instance. For example, OSPF packets can carry a different DSCP marking for the base instance and then for a VPRN service. ISIS and ARP traffic is not an IP-generated traffic type and is not configurable.

When an application is configured to use a specified DSCP value then the MPLS EXP, Dot1P bits will be marked in accordance with the network or access egress policy as it applies to the logical interface the packet will be egressing.

The DSCP value can be set per application. This setting will be forwarded to the egress line card. The egress line card does not alter the coded DSCP value and marks the LSP-EXP and IEEE 802.1p (Dot1P) bits according to the appropriate network or access QoS policy.

**Table 33** DSCP/FC Marking

| Protocol  | IPv4 | IPv6 | DSCP Marking | Dot1P Marking | Default FC |
|-----------|------|------|--------------|---------------|------------|
| ARP       | —    | —    | —            | Yes           | NC         |
| BGP       | Yes  | Yes  | Yes          | Yes           | NC         |
| BFD       | Yes  | —    | Yes          | Yes           | NC         |
| RIP       | Yes  | Yes  | Yes          | Yes           | NC         |
| PIM (SSM) | Yes  | Yes  | Yes          | Yes           | NC         |
| OSPF      | Yes  | Yes  | Yes          | Yes           | NC         |
| SMTP      | Yes  | —    | —            | —             | AF         |

**Table 33 DSCP/FC Marking (Continued)**

| Protocol                             | IPv4 | IPv6 | DSCP Marking | Dot1P Marking | Default FC |
|--------------------------------------|------|------|--------------|---------------|------------|
| IGMP/MLD                             | Yes  | Yes  | Yes          | Yes           | AF         |
| Telnet                               | Yes  | Yes  | Yes          | Yes           | AF         |
| TFTP                                 | Yes  | —    | Yes          | Yes           | AF         |
| FTP                                  | Yes  | —    | —            | —             | AF         |
| SSH (SCP)                            | Yes  | Yes  | Yes          | Yes           | AF         |
| SNMP (get, set, etc.)                | Yes  | Yes  | Yes          | Yes           | AF         |
| SNMP trap/log                        | Yes  | Yes  | Yes          | Yes           | AF         |
| syslog                               | Yes  | Yes  | Yes          | Yes           | AF         |
| OAM ping                             | Yes  | Yes  | Yes          | Yes           | AF         |
| ICMP ping                            | Yes  | Yes  | Yes          | Yes           | AF         |
| Traceroute                           | Yes  | Yes  | Yes          | Yes           | AF         |
| TACPLUS                              | Yes  | Yes  | Yes          | Yes           | AF         |
| DNS                                  | Yes  | Yes  | Yes          | Yes           | AF         |
| SNTP/NTP                             | Yes  | —    | —            | —             | AF         |
| RADIUS                               | Yes  | —    | —            | —             | AF         |
| Cflowd                               | Yes  | —    | —            | —             | AF         |
| DHCP<br>7450 ESS and<br>7750 SR only | Yes  | Yes  | Yes          | Yes           | AF         |
| Bootp                                | Yes  | —    | —            | —             | AF         |
| IPv6 Neighbor<br>Discovery           | Yes  | —    | —            | —             | NC         |

### 3.2.6.1 Default DSCP Mapping Table

```

DSCP NamedDSCP ValueDSCP ValueDSCP ValueLabel
Decimal Hexadecimal Binary
=====
Default 00x00 0b0000000be
nc1 48 0x30 0b110000h1

```

```

nc2 56 0x38 0b111000nc
ef 46 0x2e 0b101110ef
af11100x0a0b001010assured
af12120x0c0b001100assured
af13140x0e0b001110assured
af21 18 0x12 0b01001011
af22 20 0x14 0b01010011
af23220x160b01011011
af31 26 0x1a 0b01101011
af32 28 0x1c 0b01110011
af33 30 0x1d 0b01111011
af41 34 0x22 0b100010h2
af42 36 0x24 0b100100h2
af43 38 0x26 0b100110h2

default*0

```

\*The default forwarding class mapping is used for all DSCP names/values for which there is no explicit forwarding class mapping.

### 3.2.7 Configuration of TTL Propagation for VPRN Routes

This feature allows the separate configuration of TTL propagation for in transit and CPM generated IP packets, at the ingress LER within a VPRN service context. The following commands are supported:

- config router ttl-propagate vprn-local [none | vc-only | all]
- config router ttl-propagate vprn-transit [none | vc-only | all]

You can enable TTL propagation behavior separately as follows:

- for locally generated packets by CPM (vprn-local)
- for user and control packets in transit at the node (vprn-transit)

The following parameters can be specified:

- The **all** parameter enables TTL propagation from the IP header into all labels in the stack, for VPN-IPv4 and VPN-IPv6 packets forwarded in the context of all VPRN services in the system.
- The **vc-only** parameter reverts to the default behavior by which the IP TTL is propagated into the VC label but not to the transport labels in the stack. You can explicitly set the default behavior by configuring the vc-only value.



- The **none** parameter disables the propagation of the IP TTL to all labels in the stack, including the VC label. This is needed for a transparent operation of UDP traceroute in VPRN inter-AS Option B such that the ingress and egress ASBR nodes are not traced.

This command does not use a no version.

The user can override the global configuration within each VPRN instance using the following commands:

- config service vprn ttl-propagate local [inherit | none | vc-only | all]
- config service vprn ttl-propagate transit [inherit | none | vc-only | all]

The default behavior for a VPRN instance is to inherit the global configuration for the same command. You can explicitly set the default behavior by configuring the inherit value.

This command does not have a no version.

The commands do not apply when the VPRN packet is forwarded over GRE transport tunnel.

If a packet is received in a VPRN context and a lookup is done in the Global Routing Table (GRT), (when leaking to GRT is enabled for example), the behavior of the TTL propagation is governed by the LSP shortcut configuration as follows:

- when the matching route is an RSVP LSP shortcut:
  - configure router mpls shortcut-transit-ttl-propagate
- when the matching route is an LDP LSP shortcut:
  - configure router ldp shortcut-transit-ttl-propagate

When the matching route is a RFC 3107 label route or a 6PE route, It is governed by the BGP label route configuration

When a packet is received on one VPRN instance and is redirected using Policy Based Routing (PBR) to be forwarded in another VPRN instance, the TTL propagation is governed by the configuration of the outgoing VPRN instance.

Packets that are forwarded in different contexts can use different TTL propagation over the same BGP tunnel, depending on the TTL configuration of each context. An example of this might be VPRN using a BGP tunnel and an IPv4 packet forwarded over a BGP label route of the same prefix as the tunnel.

---

## 3.2.8 CE to PE Routing Protocols

The 7750 SR and 7950 XRS VPRN supports the following PE to CE routing protocols:

- BGP
- Static
- RIP
- OSPF

### 3.2.8.1 PE to PE Tunneling Mechanisms

The 7750 SR and 7950 XRS support multiple mechanisms to provide transport tunnels for the forwarding of traffic between PE routers within the 2547bis network.

The 7750 SR and 7950 XRS VPRN implementation supports the use of:

- RSVP-TE protocol to create tunnel LSPs between PE routers
- LDP protocol to create tunnel LSP's between PE routers
- GRE tunnels between PE routers

These transport tunnel mechanisms provide the flexibility of using dynamically created LSPs where the service tunnels are automatically bound (the autobind feature) and the ability to provide certain VPN services with their own transport tunnels by explicitly binding SDPs if desired. When the autobind is used, all services traverse the same LSPs and do not allow alternate tunneling mechanisms (like GRE) or the ability to craft sets of LSPs with bandwidth reservations for specific customers as is available with explicit SDPs for the service.

### 3.2.8.2 Per VRF Route Limiting

The 7750 SR and 7950 XRS allow 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.


## 3.2.9 Spoke SDPs

Distributed services use service distribution points (SDPs) to direct traffic to another router via service tunnels. SDPs are created on each participating router and then bound to a specific service. SDP can be created as either GRE or MPLS. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Services Overview Guide* for information about configuring SDPs.

This feature provides the ability to cross-connect traffic entering on a spoke SDP, used for Layer 2 services (VLLs or VPLS), on to an IES or VPRN service. From a logical point of view, the spoke SDP entering on a network port is cross-connected to the Layer 3 service as if it entered by a service SAP. The main exception to this is traffic entering the Layer 3 service by a spoke SDP is handled with network QoS policies not access QoS policies.

Figure 16 depicts traffic terminating on a specific IES or VPRN service that is identified by the *sdp-id* and VC label present in the service packet.

**Figure 16** SDP-ID and VC Label Service Identifiers



al\_0163

Refer to “VCCV BFD support for VLL, Spoke SDP Termination on IES and VPRN, and VPLS Services” in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for information about using VCCV BFD in spoke-SDP termination.



**Note:** Spoke-SDP termination of Ipipe VLLs on VPRN is not supported in System Profile B. To determine if Ipipes are currently bound to an VPRN interface, use the **show router ldp bindings services** command before configuring profile B.

### 3.2.9.1 T-LDP Status Signaling for Spoke-SDPs Terminating on IES/VPRN


T-LDP status signaling and PW active/standby signaling capabilities are supported on ipipe and epipe spoke SDPs.

Spoke SDP termination on an IES or VPRN provides the ability to cross-connect traffic entering on a spoke SDP, used for Layer 2 services (VLLs or VPLS), on to an IES or VPRN service. From a logical point of view the spoke SDP entering on a network port is cross-connected to the Layer 3 service as if it had entered using a service SAP. The main exception to this is traffic entering the Layer 3 service using a spoke SDP is handled with network QoS policies instead of access QoS policies.

When a SAP down or SDP binding down status message is received by the PE in which the Ipipe or Ethernet Spoke-SDP is terminated on an IES or VPRN interface, the interface is brought down and all associated routes are withdrawn in a similar way when the Spoke-SDP goes down locally. The same actions are taken when the standby T-LDP status message is received by the IES/VPRN PE.

This feature can be used to provide redundant connectivity to a VPRN or IES from a PE providing a VLL service, as shown in [Figure 17](#).

**Figure 17 Active/Standby VRF Using Resilient Layer 2 Circuits**



OSSG407


### 3.2.9.2 Spoke SDP Redundancy into IES/VPRN

This feature can be used to provide redundant connectivity to a VPRN or IES from a PE providing a VLL service, as shown in [Figure 17](#), using either Epipe or Ipipe spoke-SDPs. This feature is supported on the 7450 ESS and 7750 SR only.

In [Figure 17](#), PE1 terminates two spoke-SDPs that are bound to one SAP connected to CE1. PE1 chooses to forward traffic on one of the spoke SDPs (the active spoke-SDP), while blocking traffic on the other spoke-SDP (the standby spoke-SDP) in the transmit direction. PE2 and PE3 take any spoke-SDPs for which PW forwarding standby has been signaled by PE1 to an operationally down state.

The 7450 ESS, 7750 SR, and 7950 XRS routers are expected to fulfill both functions (VLL and VPRN/IES PE), while the 7705 SAR must be able to fulfill the VLL PE function. [Figure 18](#) illustrates the model for spoke-SDP redundancy into a VPRN or IES.

**Figure 18 Spoke-SDP Redundancy Model**



OSSG408

## 3.2.10 IP-VPNs

### 3.2.10.1 Using OSPF in IP-VPNs

Using OSPF as a CE to PE routing protocol allows OSPF that is currently running as the IGP routing protocol to migrate to an IP-VPN backbone without changing the IGP routing protocol, introducing BGP as the CE-PE or relying on static routes for the distribution of routes into the service providers IP-VPN. The following features are supported:


- Advertisement/redistribution of BGP-VPN routes as summary (type 3) LSAs flooded to CE neighbors of the VPRN OSPF instance. This occurs if the OSPF route type (in the OSPF route type BGP extended community attribute carried with the VPN route) is not external (or NSSA) and the locally configured domain-id matches the domain-id carried in the OSPF domain ID BGP extended community attribute carried with the VPN route.
- OSPF sham links. A sham link is a logical PE-to-PE unnumbered point-to-point interface that essentially rides over the PE-to-PE transport tunnel. A sham link can be associated with any area and can therefore appear as an intra-area link to CE routers attached to different PEs in the VPN.

### 3.2.11 IPCP Subnet Negotiation

This feature enables negotiation between Broadband Network Gateway (BNG) and customer premises equipment (CPE) so that CPE is allocated both ip-address and associated subnet.

Some CPEs use the network up-link in PPPoE mode and perform dhcp-server function for all ports on the LAN side. Instead of wasting 1 subnet for p2p uplink, CPEs use allocated subnet for LAN portion as shown in [Figure 19](#).

**Figure 19 CPEs Network Up-link Mode**



From a BNG perspective, the given PPPoE host is allocated a subnet (instead of /32) by RADIUS, external dhcp-server, or local-user-db. And locally, the host is associated with managed-route. This managed-route will be subset of the subscriber-interface subnet (on a 7450 ESS or 7750 SR), and also, subscriber-host ip-address will be from managed-route range. The negotiation between BNG and CPE allows CPE to be allocated both ip-address and associated subnet.

### 3.2.12 Cflowd for IP-VPNs

The cflowd feature allows service providers to collect IP flow data within the context of a VPRN. This data can be used to monitor types and general proportion of traffic traversing an VPRN context. This data can also be shared with the VPN customer to see the types of traffic traversing the VPN and use it for traffic engineering.


This feature should not be used for billing purposes. Existing queue counters are designed for this purpose and provide very accurate per bit accounting records.

### 3.2.13 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)*.

This feature applies to the 7450 ESS and 7750 SR only.

The first option, referred to as Option-A ([Figure 20](#)), is considered inherent in any implementation. This method uses a back-to-back connection between separate VPRN instances in each AS. As a result, each VPRN instance views the inter-AS connection as an external interface to a remote VPRN customer site. The back-to-back VRF connections between the ASBR nodes require individual sub-interfaces, one per VRF.


**Figure 20** Inter-AS Option-A: VRF-to-VRF Model

OSSG255

The second option, referred to as Option-B ([Figure 21](#)), relies heavily on the AS Boundary Routers (ASBRs) as the interface between the autonomous systems. This approach enhances the scalability of the eBGP VRF-to-VRF solution by eliminating the need for per-VPN configuration on the ASBR(s). However it requires that the ASBR(s) provide a control plan and forwarding plane connection between the autonomous systems. The ASBR(s) are connected to the PE nodes in its local autonomous system using iBGP either directly or through route reflectors. This means the ASBR(s) receive all the VPRN information and will forward these VPRN updates, VPN-IPV4, to all its EBGP peers, ASBR(s), using itself as the next-hop. It also changes the label associated with the route. This means the ASBR(s) must maintain an associate mapping of labels received and labels issued for those routes. The peer ASBR(s) will in turn forward those updates to all local IBGP peers.




**Figure 21 Inter-AS Option-B**



This form of inter-AS VPRNs performs all necessary mapping functions and the PE routers do not need to perform any additional functions than in a non-Inter-AS VPRN.

On the 7750 SR, this form of inter-AS VPRNs does not require instances of the VPRN to be created on the ASBR, as in option-A, as a result there is less management overhead. This is also the most common form of Inter-AS VPRNs used between different service providers as all routes advertised between autonomous systems can be controlled by route policies on the ASBRs by the **config>router>bgp>transport-tunnel** CLI command. The third option, referred to as Option-C (Figure 22), allows for a higher scale of VPRNs across AS boundaries but also expands the trust model between ASNs. As a result this model is typically used within a single company that may have multiple ASNs for various reasons. This model differs from Option-B, in that in Option-B all direct knowledge of the remote AS is contained and limited to the ASBR. As a result, in option-B the ASBR performs all necessary mapping functions and the PE routers do not need to perform any additional functions than in a non-Inter-AS VPRN.

**Figure 22 Option C Example**

OSSG257

With Option-C, knowledge from the remote AS is distributed throughout the local AS. This distribution allows for higher scalability but also requires all PEs and ASBRs involved in the Inter-AS VPRNs to participate in the exchange of inter-AS routing information.

In Option-C, the ASBRs distribute reachability information for remote PE'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*. Either RSVP-TE or LDP LSP can be selected to resolve next-hop for multi-hop eBGP peering by the **config>router>bgp>transport-tunnel** CLI command.

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.

### 3.2.14 VPRN Label Security at Inter-AS Boundary

This feature allows the user to enforce security at an inter-AS boundary and to configure a router, acting in a PE role and/or in an ASBR role, to accept packets of VPRN prefixes only from direct eBGP neighbors to which it advertised a VPRN label.

---

### 3.2.14.1 Feature Configuration

To use this feature, network IP interfaces that can have the feature enabled must first be identified. Participating interfaces are identified as having the untrusted state. The router supports a maximum of 15 network interfaces that can participate in this feature.

The following command is used to enable or disable this feature.

```
config>router>interface>untrusted [default-forwarding {forward | drop}]
```

Normally, the user applies the **untrusted** command to an inter-AS interface and PIP keeps track of the untrusted status of each interface. In the data path, an inter-AS interface that is flagged by PIP will cause the default forwarding to be set to the value of the **default-forwarding** keyword (**forward** or **drop**).

For backward compatibility, **default-forwarding** on the interface is set to the **forward** value. This means that labeled packets will be checked in the normal way against the table of programmed ILMs to decide if it should be dropped or forwarded in a GRT, a VRF, or a Layer 2 service context.

If the user sets the **default-forwarding** argument to the **drop** value, all labeled packets received on that interface are dropped. For details, see [Data Path Forwarding Behavior](#).

This feature sets the default behavior for an untrusted interface in the data path and for all ILMs. To allow the data path to provide an exception to the normal way of forwarding handling away from the default for VPRN ILMs, BGP must flag those ILMs to the data path.

The user enables exceptional ILM forwarding behavior, on a per-VPN-family basis, by using the following command:

```
configure>router>bgp>neighbor-trust [vpn-ipv4] [vpn-ipv6]
```

At a high level, BGP tracks each direct eBGP neighbor over an untrusted interface and to which it sent a VPRN prefix label. For each of those VPRN prefixes, BGP programs a bit map in the ILM that indicates, on a per-untrusted interface basis, whether the matching packets must be forwarded or dropped. For details, see [CPM Behavior](#).

### 3.2.14.2 CPM Behavior

This feature affects PIP behavior for management of network IP interfaces and in BGP for the resolution of BGP VPN-IPv4 and VPN-IPv6 prefixes.

The following are characteristics of CPM behavior related to PIP and the VPRN label security at inter-AS boundary feature:

- PIP manages the status of an untrusted interface based on the user configuration on the interface, as described in [Feature Configuration](#). It programs the interface record in the data path using a 4-bit untrusted interface identification number. A trusted interface has no untrusted record.
- BGP determines the status of trusted or untrusted of an eBGP neighbor by checking the untrusted record provided by PIP for the index of the interface used by the eBGP session to the neighbor.
- BGP only tracks the status of trusted or untrusted for directly connected eBGP neighbors. The neighbor address and the local address must be on the same local subnet.
- BGP includes the neighbor status of trusted or untrusted in the tribe criteria. For example, if a group consists of two untrusted eBGP neighbors and one trusted eBGP neighbor and all three neighbors have the same local-AS, neighbor-AS, and export policy, then the result is two different tribes.

As a result, if the interface status changes from trusted to untrusted or vice-versa, the eBGP neighbors on that interface will bounce.

- When the feature is enabled for a specified VPN family and BGP advertises a label for one or more resolved VPN prefixes to a group of trusted and untrusted eBGP neighbors, it creates a 16-bit map in the ILM record in which it sets the bit position corresponding to the identification number of each untrusted interface used by a eBGP session to a neighbor to which it sent the label.

A bit in the ILM record bit-map is referred to as the untrusted interface forwarding bit. The bit position corresponding to the identification number of any other untrusted interface is left clear.

For details on the data path of the ILM bit-map record, see [Data Path Forwarding Behavior](#).

- Since the same label value is advertised for prefixes in the same VRF (label per-VRF mode) and for prefixes with the same next hop (label per-next-hop mode), BGP programs the forwarding bit position in the ILM bit map for both VPN IPv4 and VPN IPv6 prefixes sharing the same label, as long as the feature is enabled for at least one of the two VPN families.

- BGP tracks, on a per-untrusted interface basis, the number of RIB-Out entries to eBGP neighbors that reference a specific VPN label. When that reference transitions from zero to a positive value or from a positive value to zero, the label for the ILM of the VPN prefix is re-downloaded to the IOM with the forwarding bit position in the ILM bit map record updated accordingly (set or unset, respectively).

This feature supports label per-VRF and label per-next-hop modes for the PE role. The feature supports label per-next-hop mode for the ASBR role.

The feature is not supported with label per-prefix mode in a PE role and is not supported in a Carrier Supporting Carrier (CSC) PE role.

### 3.2.14.3 Data Path Forwarding Behavior

ILM forwarding on a trusted interface behaves as in prior releases and is not changed. The ILM forwarding bit map is ignored and packets are forwarded normally.

ILM forwarding on an untrusted interface follows these rules.

- Only the top-most label in the label stack in a received packet is checked against the next set of rules. The top label can correspond to any one of the following applications:
  - a transport label with a pop or swap operation: static, RSVP-TE, SR-TE, LDP, SR-ISIS, SR-OSPF, or BGP-LU
  - a BGP VPRN inter-AS option B label with a swap operation when the router acts in the ASBR role for VPN routes
  - a service delimiting label for a local VRF when the router acts as a PE in a VPRN service
- The data path checks the bit position in the bit map in the ILM record, when present, that corresponds to the untrusted interface identification number in the interface record and then makes a forwarding decision to drop or forward.

A decision to forward means that a labeled packet proceeds to the regular ILM processing and its label stack is checked against the table of programmed ILMs to decide if the packet should be:

  - dropped
  - forwarded to CPM
  - forwarded as an MPLS packet
  - forwarded as an IP packet in a GRT or a VRF context
  - forwarded as a packet in a Layer 2 service context
- The following are the processing rules of the ILM:

- interface **default-forwarding=forward** and ILM bit-map not present ⇒ forward packet
  - interface **default-forwarding=forward** and interface forwarding bit position in the ILM bit-map 1 ⇒ forward packet
  - interface **default-forwarding=forward** and interface forwarding bit position in the ILM bit-map zero ⇒ drop packet
  - interface **default-forwarding=drop** and ILM bit-map not present ⇒ drop packet
  - interface **default-forwarding=drop** and interface forwarding bit position in the ILM bit-map zero ⇒ drop packet
  - interface **default-forwarding=drop** and interface forwarding bit position in the ILM bit-map 1 ⇒ forward packet
- When the eBGP neighbor is not directly connected, BGP will not track that neighbor (see [CPM Behavior](#)). In this case, the VPRN packet is received with a transport label or without a transport label if implicit-null is enabled in LDP or RSVP-TE for the transport label. Either way, the forwarding decision for the packet is solely dictated by the configuration of the **default-forwarding** value on the incoming interface.
  - If the direct eBGP neighbor sends a VPRN packet using the MPLS-over-GRE encapsulation, the data path does not check the interface forwarding bit position in the ILM bit map. In this case, the forwarding decision of the packet is solely dictated by the configuration of the **default-forwarding** value on the incoming interface.
- SR OS eBGP neighbors never use the MPLS-over-GRE encapsulation over an inter-AS link, but third party implementations might do this.

### 3.2.15 Carrier Supporting Carrier (CSC)


Carrier Supporting Carrier (CSC) is a solution for the 7750 SR and 7950 XRS that allows one service provider (the Customer Carrier) to use the IP VPN service of another service provider (the Super Carrier) for some or all of its backbone transport. RFC 4364 defines a Carrier Supporting Carrier solution for BGP/MPLS IP VPNs that uses MPLS on the interconnections between the two service providers in order to provide a scalable and secure solution.

CSC support in SR OS allows a 7750 SR or 7950 XRS to be deployed as any of the following devices shown in [Figure 23](#):

- PE1 (service provider PE)
- CSC-CE1, CSC-CE2 and CSC-CE3 (CE device from the point of view of the backbone service provider)

- CSC-PE1, CSC-PE2 and CSC-PE3 (PE device of the backbone service provider)
- ASBR1 and ASBR2 (ASBR of the backbone service provider)

**Figure 23 Carrier Supporting Carrier Reference Diagram**



### 3.2.15.1 Terminology

CE — customer premises equipment dedicated to one particular business/enterprise

PE — service provider router that connects to a CE to provide a business VPN service to the associated business/enterprise

CSC-CE — an ASBR/peering router that is connected to the CSC-PE of another service provider for purposes of using the associated CsC IP VPN service for backbone transport

CSC-PE — a PE router belonging to the backbone service provider that supports one or more CSC IP VPN services

---

### 3.2.15.2 CSC Connectivity Models

A PE router deployed by a customer service provider to provide Internet access, IP VPNs, and/or L2 VPNs may connect directly to a CSC-PE device, or it may back haul traffic within its local “site” to the CSC-CE that provides this direct connection. Here, “site” means a set of routers owned and managed by the customer service provider that can exchange traffic through means other than the CSC service. The function of the CSC service is to provide IP/MPLS reachability between isolated sites.

The CSC-CE is a “CE” from the perspective of the backbone service provider. There may be multiple CSC-CEs at a given customer service provider site and each one may connect to multiple CSC-PE devices for resiliency/multi-homing purposes.

The CSC-PE is owned and managed by the backbone service provider and provides CSC IP VPN service to connected CSC-CE devices. In many cases, the CSC-PE also supports other services, including regular business IP VPN services. A single CSC-PE may support multiple CSC IP VPN services. Each customer service provider is allocated its own VRF within the CSC-PE; VRFs maintain routing and forwarding separation and permit the use of overlapping IP addresses by different customer service providers.

A backbone service provider may not have the geographic span to connect, with reasonable cost, to every site of a customer service provider. In this case, multiple backbone service providers may coordinate to provide an inter-AS CSC service. Different inter-AS connectivity options are possible, depending on the trust relationships between the different backbone service providers.

The CSC Connectivity Models apply to the 7750 SR and 7950 XRS only.

### 3.2.15.3 CSC-PE Configuration and Operation

This section applies to CSC-PE1, CSC-PE2 and CSC-PE3 in [Carrier Supporting Carrier Reference Diagram](#).

This section applies only to the 7750 SR.

### 3.2.15.4 CSC Interface

From the point of view of the CSC-PE, the IP/MPLS interface between the CSC-PE and a CSC-CE has these characteristics:



1. The CSC interface is associated with one (and only one) VPRN service. Routes with the CSC interface as next-hop are installed only in the routing table of the associated VPRN.
2. The CSC interface supports EBGp or IBGP for exchanging labeled IPv4 routes (RFC 3107). The BGP session may be established between the interface addresses of the two routers or else between a loopback address of the CSC-PE VRF and a loopback address of the CSC-CE. In the latter case, the BGP next-hop is resolved by either a static or OSPFv2 route.
3. An MPLS packet received on a CSC interface is dropped if the top-most label was not advertised over a BGP (RFC 3107) session associated with one of the VPRN's CSC interfaces.
4. The CSC interface supports ingress QoS classification based on 802.1p or MPLS EXP. It is possible to configure a default FC and default profile for the CSC interface.
5. The CSC interface supports QoS (re)marking for egress traffic. Policies to remark 802.1p or MPLS EXP based on forwarding-class and profile are configurable per CSC interface.
6. By associating a port-based egress queue group instance with a CSC interface, the egress traffic can be scheduled/shaped with per-interface, per-forwarding-class granularity.
7. By associating a forwarding-plane based ingress queue group instance with a CSC interface, the ingress traffic can be policed to per-interface, per-forwarding-class granularity.
8. Ingress and egress statistics and accounting are available per CSC interface. The exact set of collected statistics depends on whether a queue-group is associated with the CSC interface, the traffic direction (ingress vs. egress), and the stats mode of the queue-group policers.

An Ethernet port or LAG with a CSC interface can be configured in hybrid mode or network mode. The port or LAG supports null, dot1q or qinq encapsulation. To create a CSC interface on a port or LAG in null mode, the following commands are used:

```
config>service>vprn>nw-if>port port-id
config>service>vprn>nw-if>lag lag-id
```

To create a CSC interface on a port or LAG in dot1q mode, the following commands are used:

```
config>service>vprn>nw-if>port port-id:qtag1
config>service>vprn>nw-if>lag port-id:qtag1
```

To create a CSC interface on a port or LAG in qinq mode, the following commands are used:

```
config>service>vprn>nw-if>port port-id:qtag1.qtag2
config>service>vprn>nw-if>port port-id:qtag1.*
config>service>vprn>nw-if>lag port-id:qtag1.qtag2
config>service>vprn>nw-if>lag port-id:qtag1.*
```

A CSC interface supports the same capabilities (and supports the same commands) as a base router network interface except it does not support:

- IPv6
- LDP
- RSVP
- Proxy ARP (local/remote)
- Network domain configuration
- DHCP
- Ethernet CFM
- Unnumbered interfaces

### 3.2.15.5 QoS

#### Egress

Egress traffic on a CSC interface can be shaped and scheduled by associating a port-based egress queue-group instance with the CSC interface. The steps for doing this are summarized below:

- Step 1.** Create an egress queue-group-template.
- Step 2.** Define one or more queues in the egress queue-group. For each one specify scheduling parameters such as CIR, PIR, CBS and MBS and, if using H-QoS, the parent scheduler.
- Step 3.** Apply an instance of the egress queue-group template to the network egress context of the Ethernet port with the CSC interface. When doing so, and if applicable, associate an accounting policy and/or a scheduler policy with this instance.
- Step 4.** Create a network QoS policy.
- Step 5.** In the egress part of the network QoS policy define EXP remarking rules, if necessary.
- Step 6.** In the egress part of the network QoS policy map a forwarding-class to a queue-id using the port-redirect-group command. For example:

```
config>qos>network>egress>fc$ port-redirect-group queue 5
```

- Step 7.** Apply the network QoS policy created in step 4 to the CSC interface and specify the name of the egress queue-group created in step 1 and the specific instance defined in Step 3.

### Ingress

Ingress traffic on a CSC interface can be policed by associating a forwarding-plane based ingress queue-group instance with the CSC interface. The steps for doing this are summarized below:

- Step 1.** Create an ingress queue-group-template.
- Step 2.** Define one or more policers in the ingress queue-group. For each one specify parameters such as CIR, PIR, CBS and MBS and, if using H-Pol, the parent arbiter.
- Step 3.** Apply an instance of the ingress queue-group template to the network ingress context of the forwarding plane with the CSC interface. When doing so, and if applicable, associate an accounting policy and/or a policer-control-policy with this instance.
- Step 4.** Create a network QoS policy.
- Step 5.** In the ingress part of the network QoS policy define EXP classification rules, if necessary.
- Step 6.** In the ingress part of the network QoS policy map a forwarding-class to a policer-id using the fp-redirect-group policer command. For example:  
**config>qos>network>ingress>fc\$ fp-redirect-group policer 5**
- Step 7.** Apply the network QoS policy created in step 4 to the CSC interface and specify the name of the ingress queue-group created in step 1 and the specific instance defined in step 3.

## 3.2.15.6 MPLS

BGP-3107 is used as the label distribution protocol on the CSC interface. When BGP in a CSC VPRN needs to distribute a label corresponding to a received VPN-IPv4 route, it takes the label from the global label space. The allocated label will not be re-used for any other FEC regardless of the routing instance (base router or VPRN). If a label L is advertised to the BGP peers of CSC VPRN A then a received packet with label L as the top most label is only valid if received on an interface of VPRN A, otherwise the packet is discarded.

To use BGP-3107 as the label distribution protocol on the CSC interface, add the **family label-ipv4** command to the family configuration at the instance, group, or neighbor level. This causes the capability to send and receive labeled-IPv4 routes {AFI=1, SAFI=4} to be negotiated with the CSC-CE peers.

### 3.2.15.7 CSC VPRN Service Configuration

To configure a VPRN to support CSC service, the **carrier-carrier-vpn** command must be enabled. The command fails if the VPRN service has any existing SAP or spoke SDP interfaces. A CSC interface can be added to a VPRN (using the **network-interface** command) only if the **carrier-carrier-vpn** command is enabled.

A VPRN service with the **carrier-carrier-vpn** command may be provisioned to use **auto-bind-tunnel**, configured spoke SDPs, or some combination. All SDP types are supported except for:

- GRE SDPs
- LDP over RSVP-TE tunnel SDPs

Other aspects of VPRN configuration are the same in a CSC model as in a non-CSC model.


### 3.2.16 Node Management Using VPRN

Node management can be enabled using the local interface of any VPRN service. A management VPRN is separated from other traffic using an MPLS transport tunnel. This provides IP domain separation and security for the management domain. The management domain supports IPv4 and IPv6 address families and the AAA server is connected to the same VPRN for authentication, authorization, and accounting. The SR OS allows management using a VPRN as long the management packet is destined to a local interface of the VPRN, in addition it allows configuration of the AAA servers within a VPRN; see [Figure 24](#).



**Note:** The NSP cannot discover the node using any VPRN interfaces. If node discovery is required, then GRT-LEAKING to the system IP address must be enabled. NSP can only discover the node using the system IP address in GRT.

**Figure 24 VF Network Example**



### 3.2.16.1 VPRN Management

VPRN management can be enabled by configuring the appropriate management protocol within the VPRN from the **vprn>management** context. The following protocols are supported:

- FTP
- SSH
- Telnet
- Telnetv6

By default, all protocols are disabled. When one of these protocols is enabled, that VRF becomes a management VRF.



**Note:** An SSH server can only be enabled or disabled in VPRN. All other SSH configurations such as the SSH version, key re-exchange, SSH client/server HMAC/cipher list, and so on, are configured under the system and used globally.

### 3.2.16.2 AAA Management

The **config>system>security>password>authentication-order** command sets the authentication, authorization and accounting order for the system, including VPRNs.

The **config>system>security>profile** command is used for system local user profile configuration and is used for local user authentication and authorization, including VPRNs.

Whether AAA servers are set using the **system security** command or the **aaa remote-servers** command in the VPRN, they are used as follows:

- If there is a single server configured under VPRN management AAA, then only the VPRN management AAAs are used.

For example, the **config>system>security>password>authentication-order** lists local, TAC+, and RADIUS, while the VPRN only has a RADIUS server configured, and under **system>security** both TAC+ and RADIUS are configured. In this case, if a management packet arrives on a VPRN and the destination IP matches a local interface in the VPRN, then SR OS tries local management first, and then RADIUS as configured in the VPRN. The SR OS will not try the **system>security** AAA servers since there is an AAA server configured in VPRN.

- If there is no AAA server configured under VPRN management AAA, then the system AAA servers configured in **system>security** are used.

### 3.2.17 Traffic Leaking to GRT

When a management packet arrives on a VPRN, there is a lookup for the destination IP address of the packet. If the destination IP is resolved using VPRN and the corresponding protocol is enabled under VPRN management, then the packet is extracted to CPM.

If the destination IP address is not a VRF IP and GRT leaking is enabled, a second lookup is done in the GRT FIB. If the IP belongs to a local interface in GRT and **enable-management** is enabled under **grr-leaking**, then the packet is extracted using GRT leaking to the CPM.

Traffic leaking to Global Route Table (GRT) for the 7750 SR and 7950 XRS allows service providers to offer VPRN and Internet services to their customers over a single VRF interface. This supports IPv4 and, for the 7750 SR, requires the customer VPRN interfaces to terminate on a minimum of IOM3-XP and IMM hardware. It is also supported on the 7750 SR-c12.

Packets entering a local VRF interface can have route processing results derived from the VPRN forwarding table or the GRT. The leaking and preferred lookup results are configured on a per VPRN basis. Configuration options can be general (for example, any lookup miss in the VPRN forwarding table can be resolved in the GRT), or specific (for example, specific route(s) should only be looked up in the GRT and ignored in the VPRN). In order to provide operational simplicity and improve streamlining, the CLI configuration is all contained within the context of the VPRN service.

This feature is enabled within the VPRN service context under **config>service>vprn>grt-lookup**. This is an administrative context and provides the container under which all specific commands can be entered, except for policy definition. Policy definitions remain unchanged but are referenced from this context.

The **enable-grt** command establishes the basic functionality. When it is configured, any lookup miss in the VRF table will be resolved in the GRT, if available. By itself, this only provides part of the solution. Packet forwarding within GRT must understand how to route packets back to the proper node and to the specific VPRN from which the destination exists. Destination prefixes must be leaked from the VPRN to the GRT through the use of policy. Policies are created under the **config>router>policy-options** hierarchy. By default, the number of prefixes leaked from the VPRN to the GRT is limited to five. The **export-limit** command under the **grt-lookup** hierarchy allows the service provider to override the default, or remove the limit.

When a VPRN forwarding table consists of a default route or an aggregate route, the customer may require the service provider to poke holes in those, or provide more specific route resolution in the GRT. In this case, the service provider may configure a static-route-entry and specify the GRT as the nexthop type.

The lookup result will prefer any successful lookup in the GRT that is equal to or more specific than the static route, bypassing any successful lookup in the local VPRN.

This feature and Unicast Reverse Path Forwarding (uRPF) are mutually exclusive. When a VPRN service is configured with either of these functions, the other cannot be enabled. Also, prefixes leaked from any VPRN should never conflict with prefixes leaked from any other VPRN or existing prefixes in the GRT. Prefixes should be globally unique with the service provider network and if these are propagated outside of a single providers network, they must be from the public IP space and globally unique. Network Address Translation (NAT) is not supported as part of this feature. The following type of routes will not be leaked from VPRN into the Global Routing Table (GRT):

- Aggregate routes
- Blackhole routes
- BGP VPN extranet routes

### 3.2.18 Traffic Leaking from VPRN to GRT for IPv6

This feature allows IPv6 destination lookups in two distinct routing tables and applies only to the 7750 SR and 7950 XRS. IPv6 packets within a Virtual Private Routed Network (VPRN) service is able to perform a lookup for IPv6 destination against the Global Route Table (GRT) as well as within the local VPRN.

Currently, VPRN to VPRN routing exchange is accomplished through the use of import and export policies based on Route Targets (RTs), the creation of extranets. This new feature allows the use of a single VPRN interface for both corporate VPRN routing and other services (for example, Internet) that are reachable outside the local routing context. This feature takes advantage of the dual lookup capabilities in two separate routing tables in parallel.


This feature enables IPv6 capability in addition to the existing IPv4 VPRN-to-GRT Route Leaking feature.

### 3.2.19 RIP Metric Propagation in VPRNs

When RIP is used as the PE-CE protocol for VPRNs (IP-VPNs), the RIP metric is only used by the local node running RIP with the Customer Equipment (CE). The metric is not used to or encoded into and MP-BGP path attributes exchanged between PE routers.

The RIP metric can also be used to be exchanged between PE routers so if a customer network is dual homed to separate PEs the RIP metric learned from the CE router can be used to choose the best route to the destination subnet. By using the learned RIP metric to set the BGP MED attribute, remote PEs can choose the lowest MED and in turn the PE with the lowest advertised RIP metric as the preferred egress point for the VPRN. [Figure 25](#) shows RIP metric propagation in VPRNs.

**Figure 25** RIP Metric Propagation in VPRNs





---

### 3.2.20 NTP Within a VPRN Service

Communication to external NTP clocks through VPRNs is supported in two ways: communication with external servers and peers, and communication with external clients.

Communication with external servers and peers are controlled using the same commands as used for access via base routing (see Network Time Protocol (NTP) in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide*). Communication with external clients is controlled via the VPRN routing configuration. The support for the external clients can be as unicast or broadcast service. In addition, authentication keys for external clients are configurable on a per-VPRN basis.

Only a single instance of NTP remains in the node that is time sourced to as many as five NTP servers attached to the base or management network.

The NTP show command displays NTP servers and all known clients. Because NTP is UDP-based only, no state is maintained. As a result, the **show** command output only displays when the last message from the client was received.

### 3.2.21 PTP Within a VPRN Service

The PTP within a VPRN service provides access to the PTP clock within the 7750 SR through one or more VPRN services. Only one VPRN or the base routing instance may have configured peers, but all may have discovered peers. If desired, a limit on the maximum number of dynamic peers allowed may be configured on a per routing instance basis.

For more detail on PTP see the *7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide*.

### 3.2.22 VPN Route Label Allocation

The method used for allocating a label value to an originated VPN-IP route (exported from a VPRN) depends on the configuration of the VPRN service and its VRF export policies. SR OS supports three 3 label modes:

- label per VRF
- label per next hop (LPN)

- label per prefix (LPP)

Label per VRF is the label allocation default. It is used when the label mode is configured as VRF (or not configured) and the VRF export policies do not apply an advertise-label per-prefix action. All routes exported from the VPRN with the per-VRF label have the same label value. When the PE receives a terminating MPLS packet with a per-VRF label, the label value selects the VRF context in which to perform a forwarding table lookup and this lookup determines the outgoing interface (or set of interfaces if ECMP applies).

Label per next hop is used when the exported route is not a local or aggregate route, the label mode is configured as next-hop, and the VRF export policies do not apply an advertise-label per-prefix override. It is also used when an inactive (backup path) BGP route is exported by effect of the export-inactive-bgp command if there is no advertise-label per-prefix override. All LPN-exported routes with the same primary next hop have the same per-next-hop label value. When the PE receives a terminating MPLS packet with a per-next-hop label, the label lookup selects the outgoing interface for forwarding, without any FIB lookup that might cause problems with overlapping prefixes. LPN does not support ECMP, BGP fast reroute, QPPB, or policy accounting features that might otherwise apply.

Label per-prefix is used when a qualifying IP route is exported by matching a VRF export policy action with advertise-label per-prefix. Any IPv4 or IPv6 route that is not a local route, aggregate route, BGP-VPN route, or GRT lookup static route qualifies. With LPP, every prefix is associated with its own unique label value that does not change while the route is present in the route table. When the PE receives a terminating MPLS packet with a per-prefix label value, the packet is forwarded as if the FIB lookup found only the matching prefix route and not any of the more specific prefix routes that would normally be selected. LPP supports ECMP, QPPB, and policy accounting as part of the egress forwarding decision. It does not support BGP fast reroute or BGP sticky ECMP.

The following points summarize the logic that determines the label allocation method for an exported route:

- If the IP route is LOCAL, AGGREGATE, or BGP-VPN always use the VRF label.
- If the IP route is accepted by a VRF export policy with the advertise-label per-prefix action, use LPP.
- If the IP (BGP) route is exported by the export-inactive-bgp command (VPRN best external), use LPN.
- If the IP route is exported by a VPRN configured for label-mode next-hop, use LPN.
- Else, use the per-VRF label.

### 3.2.22.1 Configuring the Service Label Mode

To change the service label mode of the VPRN for the 7750 SR, the **config>service>vprn>label-mode {vrf | next-hop}** command is used:

The default mode (if the command is not present in the VPRN configuration) is **vrf** meaning distribution of one service label for all routes of the VPRN. When a VPRN X is configured with the **label-mode next-hop** command the service label that it distributes with an IPv4 or IPv6 route that it exports depends on the type of route as summarized in [Table 34](#).

**Table 34** Service Labels Distributed in Service Label per Next-Hop Mode

| Route Type                                                                | Distributed Service Label                                                                                 |
|---------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| remote route with IP A (associated with a SAP) as resolved next-hop       | platform-wide unique label allocated to next-hop A                                                        |
| remote route with IP B (associated with a spoke SDP) as resolved next-hop | platform-wide unique label allocated to next-hop B                                                        |
| local route                                                               | platform-wide unique label allocated to VPRN X                                                            |
| aggregate route                                                           | platform-wide unique label allocated to VPRN X                                                            |
| ECMP route                                                                | platform-wide unique label allocated to next-hop A (the lowest next-hop address in the ECMP set)          |
| BGP route with a backup next-hop (BGP FRR)                                | platform-wide unique label allocated to next-hop A (the lowest next-hop address of the primary next-hops) |

A change to the label-mode of a VPRN requires the VPRN to first be shutdown.

### 3.2.22.2 Restrictions and Usage Notes

The service label per next-hop mode has the following restrictions (applies only to the 7750 SR):

- **ECMP** — The VPRN label mode should be set to **VRF** if distribution of traffic across the multiple PE-CE next-hop interfaces of an ECMP route is desired.

- Hub and spoke VPN — The VPRN label mode should not be set to next-hop if the operator does not want the hub-connected CE to be involved in the forwarding of spoke-to-spoke traffic.
- BGP next-hop indirection — BGP next-hop indirection has no benefit in service label per next-hop mode. When the resolved next-hop interface of a BGP next-hop changes all of the affected BGP routes must be re-advertised to VPRN peers with the new service label corresponding to the new resolved next-hop.
- BGP anycast — When a PE failure results in redirection of MPLS packets to the other PE in a dual-homed pair, the service label mode is forced to VRF, for example, FIB lookup will determine the next-hop even if the label mode of the VPRN is configured as next-hop.
- U-turn routing — U-turn routing is effectively disabled by service-label per next-hop.
- Carrier Supporting Carrier — The label-mode configuration of a VPRN with CSC interfaces is ignored for BGP-3107 routes learned from connected CSC-CE devices.

### 3.2.23 VPRN Support for BGP Flowspec

When a VPRN BGP instance receives an IPv4 or IPv6 flow route, and that route is valid/best, the system attempts to construct an IPv4 or IPv6 filter entry from the NLRI contents and the actions encoded in the UPDATE message. If the attempt is successful, the filter entry is added to the system-created “fSpec-*n*” IPv4 or IPv6 embedded filter, where *n* is the service-id of the VPRN. These embedded filters may be inserted into configured IPv4 and IPv6 filter policies that are applied to ingress traffic on a selected set of the VPRN’s IP interfaces. These interfaces can include SAP and spoke SDP interfaces, but not CsC network interfaces.

When flowspec rules are embedded into a user-defined filter policy, the insertion point of the rules is configurable through the **offset** parameter of the **embed-filter** command. The sum of the *ip-filter-max-size* and *offset* must not exceed the maximum filter *entry-id* range.

### 3.2.24 MPLS Entropy Label and Hash Label

The router supports both the MPLS entropy label (RFC 6790) and the Flow Aware Transport label, known as the hash label (RFC 6391). These labels allow LSR nodes in a network to load-balance labeled packets in a much more granular fashion than allowed by simply hashing on the standard label stack. See the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR MPLS Guide* for further information.

---

### 3.2.25 LSP Tagging for BGP Next-Hops or Prefixes and BGP-LU

It is possible to constrain the tunnels used by the system for resolution of BGP next-hops or prefixes and BGP labeled unicast routes using LSP administrative tags. Refer to “LSP Tagging and Auto-Bind Using Tag Information” section of the *7450 ESS, 7750 SR, 7950 XRS, and VSR MPLS Guide* for more information.

### 3.2.26 Route Leaking from the Global Route Table to VPRN Instances

The Global Route Table (GRT) to VPRN route leaking feature allows actual routes from the global route table to be exported into specific VPRN instances allowing those routes to be used for forwarding as well as being re-advertised within the VPRN context.

There are two stages for route leaking. The first stage requires the configuration of a set of **leak-export** route policies that identifies which GRT routes are subject to being exported into VPRN services. The **configure>router>leak-export** command is used to configure between one and four route policies. The GRT routes must match those policy entries that are configured with **action accept**. In addition, the **configure>router>leak-export-limit** command is used to specify the maximum number of GRT routes that can be included in the GRT leak pool.

The second stage requires the configuration of an **import-grt** policy that specifies which routes within the GRT leak pool that are leaked into the associated VPRN instances route table. The **config>service>vprn>grt-lookup>import-grt** command is used and accepts one route policy. In order for the GRT route to be leaked into the local VPRN, the route must match a policy entry with **action accept**.

If a GRT route passes both stages it is added into the VPRN route table which it to be used for IP forwarding as well as re-advertisement within other routing protocols within the VPRN context.


Both IPv4 and IPv6 routes can be leaked using this process from the GRT into one or more VPRN instances. The GRT route types that can be leaked using this process are:

- RIP, OSPF and IS-IS routes
- Direct routes
- Static routes

### 3.3 QoS on Ingress Bindings

Traffic is tunneled between VPRN service instances on different PEs over service tunnels bound to MPLS LSPs or GRE tunnels. The binding of the service tunnels to the underlying transport is achieved either automatically (using the **auto-bind-tunnel** command) or statically (using the **spoke-sdp** command; not that under the VPRN IP interface). QoS control can be applied to the service tunnels for traffic ingressing into a VPRN service; see [Figure 26](#).

**Figure 26 Ingress QoS Control on VPRN Bindings**



An ingress queue group must be configured and applied to the ingress network FP where the traffic is received for the VPRN. All traffic received on that FP for any binding in the VPRN (either automatically or statically configured) which is redirected to a policer in the FP queue group (using **fp-redirect-group** in the network QoS policy) will be controlled by that policer. As a result, the traffic from all such bindings is treated as a single entity (per forwarding class) with regard to ingress QoS control. Any **fp-redirect-group mcast-policer**, **broadcast-policer** or **unknown-policer** commands in the network QoS policy are ignored for this traffic (IP multicast traffic would use the ingress network queues or queue group related to the network interface).

Ingress classification is based on the configuration of the ingress section of the specified network QoS policy, noting that the dot1p and exp classification is based on the outer Ethernet header and MPLS label whereas the DSCP applies to the outer IP header if the tunnel encapsulation is GRE, or the DSCP in the first IP header in the payload if **ler-use-dscp** is enabled in the ingress section of the referenced network QoS policy.

Ingress bandwidth control does not take into account the outer Ethernet header, the MPLS labels/control word or GRE headers, or the FCS of the incoming frame.

The following command configures the association of the network QoS policy and the FP queue group and instance within the network ingress of a VPRN:

```
configure
 vprn
 network
 ingress
 qos <network-policy-id> fp-redirect-group <queue-group-name>
 instance <instance-id>
```

When this command is configured, it overrides the QoS applied to the related network interfaces for unicast traffic arriving on bindings in that VPRN. The IP and IPv6 criteria statements are not supported in the applied network QoS policy.

This is supported for all available transport tunnel types and is independent of the label mode (**vrf** or **next-hop**) used within the VPRN. It is also supported for Carrier-Supporting-Carrier VPRNs.

The ingress network interfaces on which the traffic is received must be on FP2- and higher-based hardware.


## 3.4 Multicast in IP-VPN Applications

This section and its subsections focuses on Multicast in IP VPN functionality. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Multicast Routing Protocols Guide* for information about multicast protocols.

Applications for this feature include enterprise customer implementing a VPRN solution for their WAN networking needs, customer applications including stock-ticker information, financial institutions for stock and other types of trading data and video delivery systems.

Implementation of multicast in IP VPNs entails the support and separation of the providers core multicast domain from the various customer multicast domains and the various customer multicast domains from each other.

**Figure 27 Multicast in IP-VPN Applications**



OSSG063



**Figure 27** depicts an example of multicast in an IP-VPN application. The provider's domain encompasses the core routers (1 through 4) and the edge routers (5 through 10). The various IP-VPN customers each have their own multicast domain, VPN-1 (CE routers 12, 13 and 16) and VPN-2 (CE Routers 11, 14, 15, 17 and 18). Multicast in this VPRN example, the VPN-1 data generated by the customer behind router 16 will be multicast only by PE 9 to PE routers 6 and 7 for delivery to CE routers 12 and 13 respectively. Data generated for VPN-2 generated by the customer behind router 15 will be forwarded by PE 8 to PE routers 5, 7 and 10 for delivery to CE routers 18, 11, 14 and 17 respectively.

The demarcation of these domains is in the PE's (routers 5 through 10). The PE router participates in both the customer multicast domain and the provider's multicast domain. The customer's CEs are limited to a multicast adjacency with the multicast instance on the PE specifically created to support that specific customer's IP-VPN. This way, customers are isolated from the provider's core multicast domain and other customer multicast domains while the provider's core routers only participate in the provider's multicast domain and are isolated from all customers' multicast domains.

The PE for a given customer's multicast domain becomes adjacent to the CE routers attached to that PE and to all other PE's that participate in the IP-VPN (or customer) multicast domain. This is achieved by the PE who encapsulates the customer multicast control data and multicast streams inside the provider's multicast packets. These encapsulated packets are forwarded only to the PE nodes that are attached to the same customer's edge routers as the originating stream and are part of the same customer VPRN. This prunes the distribution of the multicast control and data traffic to the PEs that participate in the customer's multicast domain. The Rosen draft refers to this as the default multicast domain for this multicast domain; the multicast domain is associated with a unique multicast group address within the provider's network.

### 3.4.1 Use of Data MDTs

Using the above method, all multicast data offered by a given CE is always delivered to all other CEs that are part of the same multicast. It is possible that a number of CEs do not require the delivery of a particular multicast stream because they have no downstream receivers for a specific multicast group. At low traffic volumes, the impact of this is limited. However, at high data rates this could be optimized by devising a mechanism to prune PEs from the distribution tree that although forming part of the customer multicast have no need to deliver a given multicast stream to the CE attached to them. To facilitate this optimization, the Rosen draft specifies the use of data MDTs. These data MDTs are signaled once the bandwidth for a given SG exceeds the configurable threshold.

Once a PE detects it is transmitting data for the SG in excess of this threshold, it sends an MDT join TLV (at 60 second intervals) over the default MDT to all PEs. All PEs that require the SG specified in the MDT join TLV will join the data MDT that will be used by the transmitting PE to send the given SG. PEs that do not require the SG will not join the data MDT, thus pruning the multicast distribution tree to just the PEs requiring the SG. After providing sufficient time for all PEs to join the data MDT, the transmitting PE switches the given multicast stream to the data MDT.

PEs that do not require the SG to be delivered, keep state to allow them to join the data MDT as required.

When the bandwidth requirement no longer exceeds the threshold, the PE stops announcing the MDT join TLV. At this point the PEs using the data MDT will leave this group and transmission resumes over the default MDT.

Sampling to check if an s,g has exceeded the threshold occurs every ten seconds. If the rate has exceeded the configured rate in that sample period then the data MDT is created. If during that period the transmission rate has not exceeded the configured threshold then the data MDT is not created. If the data MDT is active and the transmission rate in the last sample period has not exceeded the configured rate then the data MDT is torn down and the multicast stream resumes transmission over the default MDT.

### 3.4.2 Multicast Protocols Supported in the Provider Network

When MVPN auto-discovery is disabled, PIM-SM can be used for I-PMSI, and PIM-SSM or PIM-SM (Draft-Rosen Data MDT) can be used for S-PMSI; When MVPN S-PMSI auto-discovery is enabled, both PIM-SM and PIM SSM can be used for I-PMSI, and PIM-SSM can be used for S-PMSI. In the customer network, both PIM-SM and PIM-SSM are supported.

An MVPN is defined by two sets of sites: sender sites set and receiver sites set, with the following properties:

- Hosts within the sender sites set could originate multicast traffic for receivers in the receiver sites set.
- Receivers not in the receiver sites set should not be able to receive this traffic.
- Hosts within the receiver sites set could receive multicast traffic originated by any host in the sender sites set.
- Hosts within the receiver sites set should not be able to receive multicast traffic originated by any host that is not in the sender sites set.

A site could be both in the sender sites set and receiver sites set, which implies that hosts within such a site could both originate and receive multicast traffic. An extreme case is when the sender sites set is the same as the receiver sites set, in which case all sites could originate and receive multicast traffic from each other.

Sites within a given MVPN may be either within the same, or in different organizations, which implies that an MVPN can be either an intranet or an extranet. A given site may be in more than one MVPN, which implies that MVPNs may overlap. Not all sites of a given MVPN have to be connected to the same service provider, which implies that an MVPN can span multiple service providers.

Another way to look at MVPN is to say that an MVPN is defined by a set of administrative policies. Such policies determine both sender sites set and receiver site set. Such policies are established by MVPN customers, but implemented by MVPN service providers using the existing BGP/MPLS VPN mechanisms, such as route targets, with extensions, as necessary.

### 3.4.3 MVPN Membership Auto-discovery using BGP

BGP-based auto-discovery is performed by a multicast VPN address family. Any PE that attaches to an MVPN must issue a BGP update message containing an NLRI in this address family, along with a specific set of attributes.

The PE router uses route targets to specify MVPN route import and export. The route target may be the same as the one used for the corresponding unicast VPN, or it may be different. The PE router can specify separate import route targets for sender sites and receiver sites for a given MVPN.

The route distinguisher (RD) that is used for the corresponding unicast VPN can also be used for the MVPN.

When BGP auto-discovery is enabled, PIM peering on the I-PMSI is disabled, so no PIM hellos are sent on the I-PMSI. C-trees to P-tunnels bindings are also discovered using BGP S-PMSI AD routes, instead of PIM join TLVs. Configure PIM join TLVs when **c-mcast-signaling** is set to **pim** in the **config>service>vprn>mvpn>provider-tunnel>selective>auto-discovery-disable** context.

[Table 35](#) and [Table 36](#) describe the supported configuration combinations. If the CLI combination is not allowed, the system returns an error message. If the CLI command is marked as “ignored” in the table, the configuration is not blocked, but its value is ignored by the software.

**Table 35** Supported Configuration Combinations

| Auto-Discovery | Inclusive PIM SSM  | Action      |
|----------------|--------------------|-------------|
| Yes            | Yes                | Allowed     |
| MDT-SAFI       | Yes                | Allowed     |
| No             | Yes                | Not Allowed |
| Yes or No      | No                 | Allowed     |
| MDT-SAFI       | No                 | Ignored     |
| MDT-SAFI       | No (RSVP and MLDP) | Not Allowed |

**Table 36** Supported Configuration Combinations

| Auto-Discovery | C-Mcast-Signaling      | s-PMSI Auto-Discovery   | Action      |
|----------------|------------------------|-------------------------|-------------|
| Yes            | BGP                    | Ignored                 | Allowed     |
| Yes            | PIM                    | Yes                     | Allowed     |
| Yes            | PIM                    | No                      | Allowed     |
| No             | BGP                    | Ignored                 | Not Allowed |
| No             | PIM                    | Ignored                 | Allowed     |
| MDT-SAFI       | Ignored (PIM behavior) | Ignored ("No" behavior) | Allowed     |

For example, if **auto-discovery** is disabled, the **c-mcast-signaling bgp** command will fail with an error message stating:

C-multicast signaling in BGP requires auto-discovery to be enabled

If **c-mcast-signaling** is set to **bgp** then **no auto-discovery** will fail with an error message stating

C-multicast signaling in BGP requires auto-discovery to be enabled

When **c-mcast-signaling** is set to **bgp**, S-PMSI A-D is always enabled (its configuration is ignored);

When **auto-discovery** is disabled, S-PMSI A-D is always disabled (its configuration is ignored).

When **auto-discovery** is enabled and **c-multicast-signaling** is set to **pim**, the S-PMSI A-D configuration value is used.

**mdt-safi** uses **pim c-mcast-signaling** and **s-pmsi-signaling** regardless of what is configured. A **c-mcast-signaling** or **s-pmsi-signaling** configuration is ignored, but both **pim** and **bgp** values are allowed.

**mdt-safi** is only applicable to PIM-SSM I-PMSI. PIM-SM (ASM) I-PMSI is configurable but is ignored. RSVP and MLDP I-PMSI are not allowed.

MVPN implementation based on RFC 6037, *Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs* can support membership auto-discovery using BGP MDT-SAFI. A CLI option is provided per MVPN instance to enable auto-discovery either using BGP MDT-SAFI or NG-MVPN. Only PIM-MDT is supported with BGP MDT-SAFI method.

### 3.4.4 PE-PE Transmission of C-Multicast Routing using BGP

MVPN c-multicast routing information is exchanged between PEs by using c-multicast routes that are carried using MCAST-VPN NLRI.

### 3.4.5 VRF Route Import Extended Community

VRF route import is an IP address-specific extended community, of an extended type, and is transitive across AS boundaries (RFC 4360, *BGP Extended Communities Attribute*).

To support MVPN, in addition to the import/export route target extended communities used by the unicast routing, each VRF on a PE must have an import route target extended community that controls imports of C-multicast routes into a particular VRF.

The c-multicast import RT uniquely identifies a VRF, and is constructed as follows:

- The Global Administrator field of the c-multicast import RT must be set to an IP address of the PE. This address should be common for all the VRFs on the PE (this address may be the PE's loopback address).
- The Local Administrator field of the c-multicast import RT associated with a given VRF contains a 2 octets long number that uniquely identifies that VRF within the PE that contains the VRF.

A PE that has sites of a given MVPN connected to it communicates the value of the c-multicast import RT associated with the VRF of that MVPN on the PE to all other PEs that have sites of that MVPN. To accomplish this, a PE that originates a (unicast) route to VPN-IP addresses includes in the BGP updates message that carries this route the VRF route import extended community that has the value of the c-multicast import RT of the VRF associated with the route, except if it is known a priori that none of these addresses will act as multicast sources and/or RP, in which case the (unicast) route need not carry the VRF Route Import extended community.

All c-multicast routes with the c-multicast import RT specific to the VRF must be accepted. In this release, vrf-import and vrf-target policies don't apply to C-multicast routes.

The decision flow path is shown below.

```
if (route-type == c-mcast-route)
 if (route_target_list includes C-multicast_Import_RT){
 else
 drop;
 else
 Run vrf_import and/or vrf-target;
```

## 3.4.6 Provider Tunnel Support

### 3.4.6.1 Point-to-Multipoint Inclusive (I-PMSI) and Selective (S-PMSI) Provider Multicast Service Interface

BGP C-multicast signaling must be enabled for an MVPN instance to use P2MP RSVP-TE or LDP as I-PMSI (equivalent to 'Default MDT', as defined in draft Rosen MVPN) and S-PMSI (equivalent to 'Data MDT', as defined in draft Rosen MVPN).

By default, all PE nodes participating in an MVPN receive data traffic over I-PMSI. Optionally, (C-\*, C-\*) wildcard S-PMSI can be used instead of I-PMSI. See section [Wildcard \(C-\\*, C-\\*\) P2MP LSP S-PMSI](#) for more information. For efficient data traffic distribution, one or more S-PMSIs can be used, in addition to the default PMSI, to send traffic to PE nodes that have at least one active receiver connected to them. For more information, see [P2MP LSP S-PMSI](#).

Only one unique multicast flow is supported over each P2MP RSVP-TE or P2MP LDP LSP S-PMSI. Number of S-PMSI that can be initiated per MVPN instance is restricted by CLI command **maximum-p2mp-spmsi**. P2MP LSP S-PMSI cannot be used for more than one (S,G) stream (that is, multiple multicast flow) as number of S-PMSI per MVPN limit is reached. Multicast flows that cannot switch to S-PMSI remain on I-PMSI.

### 3.4.6.2 P2MP RSVP-TE I-PMSI and S-PMSI

Point-to-Multipoint RSVP-TE LSP as inclusive or selective provider tunnel is available with BGP NG-MVPN only. P2MP RSVP-TE LSP is dynamically setup from root node on auto discovery of leaf PE nodes that are participating in multicast VPN. Each RSVP-TE I-PMSI or S-PMSI LSP can be used with a single MVPN instance only.

RSVP-TE LSP template must be defined (see *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR MPLS Guide*) and bound to MVPN as inclusive or selective (S-PMSI is for efficient data distribution and is optional) provider tunnel to dynamically initiate P2MP LSP to the leaf PE nodes learned via NG-MVPN auto-discovery signaling. Each P2MP LSP S2L is signaled based on parameters defined in LSP template.

### 3.4.6.3 P2MP LDP I-PMSI and S-PMSI

Point-to-Multipoint LDP LSP as inclusive or selective provider tunnel is available with BGP NG-MVPN only. P2MP LDP LSP is dynamically setup from leaf nodes on auto discovery of leaf node PE nodes that are participating in multicast VPN. Each LDP I-PMSI or S-PMSI LSP can be used with a single MVPN instance only.

The **multicast-traffic** CLI command must be configured per LDP interface to enable P2MP LDP setup. P2MP LDP must also be configured as inclusive or selective (S-PMSI is for efficient data distribution and is optional) provider tunnel per MVPN to dynamically initiate P2MP LSP to leaf PE nodes learned via NG-MVPN auto-discovery signaling.

### 3.4.6.4 Wildcard (C-\*, C-\*) P2MP LSP S-PMSI

Wildcard S-PMSI allows usage of selective tunnel as a default tunnel for a given MVPN. By using wildcard S-PMSI, operators can avoid full mesh of LSPs between MVPN PEs, reducing related signaling, state, and BW consumption for multicast distribution (no traffic is sent to PEs without any receivers active on the default PMSI).

The SR OS allows an operator to configure wildcard S-PMSI for ng-MVPN (**config>service>vprn>mvpn>pt>inclusive>wildcard-spmsi**), using LDP and RSVP-TE in P-instance. Support includes:

- IPv4 and IPv6
- PIM ASM and SSM
- directly attached receivers

The SR OS (C-\*, C-\*) wildcard implementation uses wildcard S-PMSI instead of I-PMSI for a given MVPN. To switch MVPN from I-PMSI to (C-\*, C-\*) S-PMSI a VPRN shutdown is required. ISSU and UMH redundancy can be used to minimize the impact.

To minimize outage, the following upgrade order is recommended:

1. Route Reflector
2. Receiver PEs
3. backup UMH
4. active UMH

RSVP-TE/mLDP configuration under inclusive provider tunnel (**config>service>vprn>mvpn>pt>inclusive**) apply to wildcard S-PMSI when enabled.

Wildcard C-S and C-G values are encoded as defined in RFC6625: using zero for Multicast Source Length and Multicast Group Length and omitting Multicast Source and Multicast Group values respectively in MCAST\_VPN\_NLRI. For example, a (C-\*, C-\*) will be advertised as: RD, 0x00, 0x00, and originating router's IP address.





**Note:** All SR OSs with BGP peering session to the PE with RFC6625 support enabled must be upgraded to SR OS release 13.0 before the feature is enabled. Failure to do so will result in the following processing on a router with BGP peering session to an RFC6625-enabled PE:

- BGP peer running Release 12.0 version R4 or newer will accept 0-length address and it will keep encoding length 4 with all zeros for the address
- BGP peer running Release 12 version R3 or older will not accept 0-length address and will keep restarting BGP session

The procedures implemented by SR OS are compliant to section 3 and 4 of RFC, 6625. Wildcards encoded as described above are carried in NLRI field of MP\_REACH\_NLRF\_ATTRIBUTE. Both IPv4 and IPv6 are supported: (AFI) of 1 or 2 and a Subsequent AFI (SAFI) of MCAST-VPN.

The (C-\*, C-\*) S-PMSI is established as follows:

- UMH PEs advertise I-PMSI A-D routes without tunnel information present (empty PTA) - encoded as per RFC6513/6514 prior to advertising wildcard S-PMSI. I-PMSI needs to be signaled and installed on receiver PEs, because (C-\*, C-\*) S-PMSI is only installed when a first receiver is added. However, no LSP is established for I-PMSI).
- UMH PEs advertise S-PMSI A-D route whose NLRI contains (C-\*, C-\*) with tunnel information encoded as per RFC 6625.
- Receiver PEs join wildcard S-PMSI if there are any receivers present.



**Note:** If UMH PE does not encode I-PMSI/S-PMSI A-D routes as per the above, or advertises both I-PMSI and wildcard S-PMSI with the tunnel information present, no interoperability can be achieved.

To ensure proper operation of BSR between PEs with (C-\*, C-\*) S-PMSI signaling, SR OS implements two modes of operations for BSR.

By default (bsr unicast):

- BSR PDUs are sent/forwarded as unicast PDUs to neighbor PEs when I-PMSI with Pseudo-tunnel interface is installed.
- At every BSR interval timer BSR Unicast PDU are sent to all I-PMSI interfaces when this is an elected BSR.
- BSMs received as multicast from C-instance interfaces are flooded as unicast in the P-instance.
- All PEs process BSR PDU's received on I-PMSI Pseudo-tunnel interface as unicast packets.

- BSR PDU's are not forwarded to PE's management control interface.
- BSR unicast PDU's use PE's System IP address as destination IP and sender PE's System address as Source IP.
- The BSR unicast functionality ensures that no special state needs to be created for BSR when (C-\*, C-\*) S-PMSI is enabled, which is beneficiary considering low volume of BSR traffic.

**Note:**

- For bsr unicast, the base IPv4 system address (IPv4) or the mapped version of the base IPv4 system address (IPv6) must be configured under the VPRN to ensure bsr unicast messages can reach the VPRN.
- For bsr spmsi, the base IPv4/IPv6 system address must be configured under the VPRN to ensure B-SR S-PMSI's are established.

BSR S-PMSI mode can be enabled to allow interoperability with other vendors. In that mode full mesh S-PMSI is required and created between all PEs in MVPN to exchange BSR PDUs between all PEs in MVPN. To operate as expected, the BSR S-PMSI mode requires a selective P-tunnel configuration. For IPv6 support (including dual-stack) of BSR S-PMSI mode, the IPv6 default system interface address must be configured as a loopback interface address under the VPRN and VPRN PIM contexts. Changing BSR signaling requires a VPRN shutdown.

Other key feature interactions and caveats for (C-\*, C-\*) include the following:

- Extranet is fully supported with wildcard S-PMSI trees.
- (C-S, C-G) S-PMSIs are supported when (C-\*, C-\*) S-PMSI is configured (including both BW and receiver PE driven thresholds).
- Geo-redundancy is supported (deploying with geo-redundancy eliminates traffic duplication when geo-redundant source has no active receivers at a cost of slightly increased outage upon a switch since wildcard S-PMSI may need to be re-establish).
- PIM in P-instance is not supported.
- SR OS implementation requires wildcard encoding as per RFC6625 and I-PMSI/ S-PMSI signaling as defined above (I-PMSI signaled with empty PTA then S-PMSI signaled with P-tunnel PTA) for interoperability. Implementations that do not adhere to RFC6625 encoding, or signal both I-PMSI and S-PMSI with P-tunnel PTA will not inter-operate with SR OS implementation).

---

### 3.4.6.5 P2MP LSP S-PMSI

NG-MVPN support P2MP RSVP-TE and P2MP LDP LSPs as selective provider multicast service interface (S-PMSI). S-PMSI is used to avoid sending traffic to PEs that participate in multicast VPN, but do not have any receivers for a given C-multicast flow. This allows more-BW efficient distribution of multicast traffic over the provider network, especially for high bandwidth multicast flows. S-PMSI is spawned dynamically based on configured triggers as described in S-PMSI trigger thresholds section.

In MVPN, the head-end PE firstly discovers all the leaf PEs via I-PMSI A-D routes. It then signals the P2MP LSP to all the leaf PEs using RSVP-TE. In the scenario of S-PMSI:

1. The head-end PE sends an S-PMSI A-D route for a specific C-flow with the Leaf Information Required bit set.
2. The PEs who are interested in the C-flow respond with Leaf A-D routes.
3. The head-end PE then signals the P2MP LSP to all the leaf PEs using RSVP-TE.

Also, because the receivers may come and go, the implementation supports dynamically adding and pruning leaf nodes to and from the P2MP LSP.

When the tunnel type in the PMSI attribute is set to RSVP-TE P2MP LSP, the tunnel identifier is <Extended Tunnel ID, Reserved, Tunnel ID, P2MP ID>, as carried in the RSVP-TE P2MP LSP SESSION Object.

The PE can also learn via an A-D route that it needs to receive traffic on a particular RSVP-TE P2MP LSP before the LSP is actually setup. In this case, the PE needs to wait until the LSP is operational before it can modify its forwarding tables as directed by the A-D route.

Because of the way that LDP normally works, mLDP P2MP LSPs are setup without solicitation from the leaf PEs towards the head-end PE. The leaf PE discovers the head-end PE via I-PMSI or S-PMSI A-D routes. The tunnel identifier carried in the PMSI attribute is used as the P2MP FEC element. The tunnel identifier consists of the head-end PE's address, along with a Generic LSP identifier value. The Generic LSP identifier value is automatically generated by the head-end PE.


### 3.4.6.6 Dynamic Multicast Signaling over P2MP LDP in VRF

This feature provides a multicast signaling solution for IP-VPNs, allowing the connection of IP multicast sources and receivers in C-instances, which are running PIM multicast protocol using Rosen MVPN with BGP SAFI and P2MP mLDP in P-instance. The solution dynamically maps each PIM multicast flow to a P2MP LDP LSP on the source and receiver PEs.

The feature uses procedures defined in RFC 7246: *Multipoint Label Distribution Protocol In-Band Signaling in Virtual Routing and Forwarding (VRF) Table Context*. On the receiver PE, PIM signaling is dynamically mapped to the P2MP LDP tree setup. On the source PE, signaling is handed back from the P2MP mLDP to the PIM. Due to dynamic mapping of multicast IP flow to P2MP LSP, provisioning and maintenance overhead is eliminated as multicast distribution services are added and removed from the VRF. Per (C-S, C-G) IP multicast state is also removed from the network, since P2MP LSPs are used to transport multicast flows.

Figure 28 illustrates dynamic mLDP signaling for IP multicast in VPRN.

**Figure 28** Dynamic mLDP Signaling for IP Multicast in VPRN



al\_0619

As illustrated in Figure 28, P2MP LDP LSP signaling is initiated from the receiver PE that receives PIM JOIN from a downstream router (Router A). To enable dynamic multicast signaling, the **p2mp-ldp-tree-join** must be configured on PIM customer-facing interfaces for the given VPRN of Router A. This enables handover of multicast tree signaling from the PIM to the P2MP LDP LSP. Being a leaf node of the P2MP LDP LSP, Router A selects the upstream-hop as the root node of P2MP LDP FEC,

based on a routing table lookup. If an ECMP path is available for a given route, then the number of trees are equally balanced towards multiple root nodes. The PIM joins are carried in the Transit Source PE (Router B), and multicast tree signaling is handed back to the PIM and propagated upstream as native-IP PIM JOIN toward C-instance multicast source.

The feature is supported with IPv4 and IPv6 PIM SSM and IPv4 mLDP. Directly connected IGMP/MLD receivers are also supported, with PIM enabled on outgoing interfaces and SSM mapping configured, if required.

The following are feature caveats:

- Dynamic mLDP signaling in a VPRN instance is mutually exclusive with Rosen or NG-MVPN.
- A single instance of P2MP LDP LSP is supported between the receiver PE and Source PE per multicast flow; there is no stitching of dynamic trees.
- Extranet functionality is not supported.
- The router LSA link ID or the advertising router ID must be a routable IPv4 address (including IPv6 into IPv4 mLDP use cases).
- IPv6 PIM with dynamic IPv4 mLDP signaling is not supported with e-BGP or i-BGP with IPv6 next-hop.
- Inter-AS and IGP inter-area scenarios where the originating router is altered at the ASBR and ABR respectively, (hence PIM has no way to create the LDP LSP towards the source), are not supported.
- When dynamic mLDP signaling is deployed, a change in Route Distinguisher (RD) on the Source PE is not acted upon for any (C-S, C-G)s until the receiver PEs learn about the new RD (via BGP) and send explicit delete and create with the new RD.
- Procedures of Section 2 of RFC 7246 for a case where UMH and the upstream PE do not have the same IP address are not supported.

### **3.4.6.7 MVPN Sender-only/Receiver-only**


In multicast MVPN, by default, if multiple PE nodes form a peering with a common MVPN instance then each PE node originates a multicast tree locally towards the remaining PE nodes that are member of this MVPN instance. This behavior creates a mesh of I-PMSI across all PE nodes in the MVPN. It is often a case, that a given VPN has many sites that will host multicast receivers, but only few sites that either host both receivers and sources or sources only.

MVPN Sender-only/Receiver-only allows optimization of control and data plane resources by preventing unnecessary I-PMSI mesh when a given PE hosts multicast sources only or multicast receivers only for a given MVPN.

For PE nodes that host only multicast sources for a given VPN, operators can now block those PEs, through configuration, from joining I-PMSIs from other PEs in this MVPN. For PE nodes that host only multicast receivers for a given VPN, operators can now block those PEs, through configuration, to set-up a local I-PMSI to other PEs in this MVPN.

MVPN Sender-only/Receiver-only is supported with ng-MVPN using IPv4 RSVP-TE or IPv4 LDP provider tunnels for both IPv4 and IPv6 customer multicast. [Figure 29](#) depicts 4-site MVPN with sender-only, receiver-only and sender-receiver (default) sites.

**Figure 29 MVPN Sender-only/Receiver-only Example**



Extra attention needs to be paid to BSR/RP placement when Sender-only/Receiver-only is enabled. Source DR sends unicast encapsulated traffic towards RP, therefore, RP shall be at sender-receiver or sender-only site, so that \*G traffic can be sent over the tunnel. BSR shall be deployed at the sender-receiver site. BSR can be at sender-only site if the RPs are at the same site. BSR needs to receive packets from other candidate-BSR and candidate-RPs and also needs to send BSM packets to everyone.

### 3.4.6.8 S-PMSI Trigger Thresholds

The mLDP and RSVP-TE S-PMSIs support two types of data thresholds: bandwidth-driven and receiver-PE-driven. The threshold evaluation and bandwidth driven threshold functionality are described in [Use of Data MDTs](#).

In addition to the bandwidth threshold functionality, an operator can enable receiver-PE-driven threshold behavior. Receiver PE thresholds ensure that S-PMSI is only created when BW savings in P-instance justify extra signaling required to establish a new S-PMSI. For example, the number of receiver PEs interested in a given C-multicast flow is meaningfully smaller than the number of receiver PEs for default PMSI (I-PMSI or wildcard S-PMSI). To ensure that S-PMSI is not constantly created/deleted, two thresholds need to be specified: receiver PE add threshold and receiver PE delete threshold (expected to be significantly higher).

When a (C-S, C-G) crosses a data threshold to create S-PMSI, instead of regular S-PMSI signaling, sender PE originates S-PMSI explicit tracking procedures to detect how many receiver PEs are interested in a given (C-S, C-G). When receiver PEs receive an explicit tracking request, each receiver PE responds, indicating whether there are multicast receivers present for that (C-S, C-G) on the given PE (PE is interested in a given (C-S, C-G)). If the geo-redundancy feature is enabled, receiver PEs do not respond to explicit tracking requests for suppressed sources and therefore only Receiver PEs with an active join are counted against the configured thresholds on Source PEs.

Upon regular sampling and check interval, if the previous check interval had a non-zero receiver PE count (one interval delay to not trigger S-PMSI prematurely) and current count of receiver PEs interested in the given (C-S, C-G) is non-zero and is less than the configured receiver PE add threshold, Source PE will set-up S-PMSI for this (C-S, C-G) following standard ng-MVPN procedures augmented with explicit tracking for S-PMSI being established.

Data threshold timer should be set to ensure enough time is given for explicit tracking to complete (for example, setting the timer to value that is too low may create S-PMSI prematurely).

Upon regular data-delay-interval expiry processing, when BW threshold validity is being checked, a current receiver PE count is also checked (for example, explicit tracking continues on the established S-PMSI). If BW threshold no longer applies or the receiver PEs exceed receiver PE delete threshold, the S-PMSI is torn down and (C-S, C-G) joins back the default PMSI.

Changing of thresholds (including enabling disabling the thresholds) is allowed in service. The configuration change is evaluated at the next periodic threshold evaluation.

The explicit tracking procedures follow RFC6513/6514 with clarification and wildcard S-PMSI explicit tracking extensions as described in IETF Draft: *draft-dolganow-l3vpn-expl-track-00*.

### 3.4.6.9 Migration from Existing Rosen Implementation

The existing Rosen implementation is compatible to provide an easy migration path.

The following migration procedure are supported:

- Upgrade all the PE nodes that need to support MVPN to the newer release.
- The old configuration will be converted automatically to the new style.
- Node by node, MCAST-VPN address-family for BGP is enabled. Enable auto-discovery using BGP.
- Change PE-to-PE signaling to BGP.

### 3.4.6.10 Policy-based S-PMSI

SR OS creates a single Selective P-Multicast Service Interface (S-PMSI) per multicast stream: (S,G) or (\*,G). To better manage bandwidth allocation in the network, multiple multicast streams are often bundled starting from the same root node into a single, multi-stream S-PMSI. Network bandwidth is usually managed per package or group of packages, rather than per channel.

Multi-stream S-PMSI supports a single S-PMSI for one or more IPv4 (C-S, C-G) or IPv6 (C-S, C-G) streams. Multiple multicast streams starting from the same VPRN going to the same set of leafs can use a single S-PMSI. Multi-stream S-PMSIs can:

- carry exclusively IPv4 or exclusively IPv6, or a mix of channels
- co-exist with a single group S-PMSI



To create a multi-stream S-PMSI, an S-PMSI policy needs to be configured in the VPN context on the source node. This policy maps multiple (C-S, C-G) streams to a single S-PMSI. Because this configuration is done per MVPN, multiple VPNs can have identical policies, each configured for its own VPN context.

When mapping a multicast stream to a multi-stream S-PMSI policy, the data will traverse the S-PMSI without first using the I-PMSI. (Before this feature, when a multicast stream was sourced, the data used the I-PMSI first until a configured threshold was met. After this, if the multicast data exceeded the threshold, it would signal the S-PMSI tunnel and switch from I-PMSI to S-PMSI.)

For multi-stream S-PMSI, if the policy is configured and the multicast data matches the policy rules, the multi-stream S-PMSI is used from the start without using the default I-PMSI.

Multiple multi-stream S-PMSI policies could be assigned to a specific S-PMSI configuration. In this case, the policy acts as a link list. The first (lowest index) that matches the multi-stream S-PMSI policy is used for that specific stream.

The rules for matching a multi-stream S-PMSI on the source node are listed here.

S-PMSI to (C-S, C-G) mapping on Source-PE, in sequence:

1. The multi-stream S-PMSI policy is evaluated, starting from the lowest numerical policy index. This allows the feature to be enabled in the service when per-(C-S, C-G) stream configuration is present. Only entries that are not shut down are evaluated. First, the multi-stream S-PMSI (the lowest policy index) that the (C-S, C-G) stream maps to is selected.
2. If (C-S, C-G) does not map to any of multi-stream S-PMSIs, per-(C-S, C-G) S-PMSIs are used for transmission if a (C-S, C-G) maps to an existing S-PMSI (based on data-thresholds).
3. If S-PMSI cannot be used, the default PMSI is used.

Multi-stream S-PMSI P-tunnel failure handling:

- If an S-PMSI P-tunnel is not available, a default PMSI tunnel is used. When an S-PMSI tunnel fails, all (C-S, C-G) streams using this multi-stream S-PMSI move to the default PMSI. The groups move back to S-PMSI after the S-PMSI tunnel is restored.

### 3.4.6.10.1 Supported MPLS Tunnels

Multi-stream S-PMSI is configured in the context of an **auto-discovery** default (that is, NG-MVPN). It supports all existing per-mLDP/RSVP-TE P2MP S-PMSI tunnel functionality for multi-stream S-PMSI LSP templates (RSVP-TE P-instance).

Notes:

- Per-multicast group statistics are not available for multi-stream S-PMSIs on S-PMSI level.
- GRE tunnels are not supported for multi-stream S-PMSI.

### 3.4.6.10.2 Supported Multicast Features

S-PMSI is supported with PIM ASM and PIM SSM in C-instances.

Notes:

- The multi-stream S-PMSI model uses BSR RP co-located with the source PE or an RP between the source PE and multicast source, that is, upstream of receivers. Both **bsr unicast** and **bsr spmsi** can be deployed as applicable.
- The model also supports other RP types.

### 3.4.6.10.3 In-service Changes to Multi-stream S-PMSI

The operator can change the mapping in service; that is, the operator can move active streams (C-S, C-G) from one S-PMSI to another using the configuration, or from the default PMSI to the S-PMSI, without having to stop data transmission or having to disable a PMSI.

The change is performed by moving a (C-S, C-G) stream from a per-group S-PMSI to a multi-stream S-PMSI and vice versa, and moving a (C-S, C-G) stream from one multi-stream S-PMSI to another multi-stream S-PMSI.

Notes:

- During re-mapping, a changed (C-S, C-G) stream will first be moved to the default PMSI before it is moved to a new S-PMSI, regardless of the type of move. Unchanged (C-S, C-G) streams must remain on an existing PMSIs.

- Any change to a multi-stream S-PMSI policy or to a preferred multi-stream S-PMSI policy (for example, an index change equivalent to less than or equal to the current policy) should be performed during a maintenance window. Failure to perform these types of changes in a maintenance window could potentially cause a traffic outage.

#### 3.4.6.10.4 Configuration Example

In this example, two policies are created on the source node: multi-stream S-PMSI 1 and multi-stream-S-PMSI 10.

A multicast stream with group 224.0.0.x and source 192.0.2.0/24 will map to the first multi-stream policy. The group in the range of 224.0.0.0/24 and source 192.0.2.1/24 will map to policy 10.

```
*A:SwSim14>config>service>vprn# info
mvpn
 auto-discovery default
 c-mcast-signaling bgp
 provider-tunnel
 inclusive
 mldp
 no shutdown
 exit
 exit
 selective
 mldp
 no shutdown
 exit
 no auto-discovery-disable
 data-threshold 239.70.1.0/24 1
 data-threshold 239.0.0.0/8 1
 multistream-spmsi 1 create
 group 224.0.0.0/24
 source 192.0.2.0/24
 exit
 exit
 multistream-spmsi 10 create
 group 224.0.0.0/24
 source 192.0.2.1/24
 exit
 exit
 exit
exit
```

### 3.4.6.11 Policy-based Data MDT

A single data MDT can transport one or more IPv4 (C-S, C-G) streams. This allows multiple multicast streams starting from the same VPRN going to the same set of leafs to utilize a single data MDT. Characteristics of an MDT include:

- a multi-stream data MDT can carry IPv4 only
- a multi-stream data MDT can co-exist with a single-group data MDT
- a default MDT must be configured

To create a multi-stream MDT data, an MDA data policy must be configured in the context of MVPN on the source node. This policy maps multiple (C-S, C-G) streams to a single data MDT. Because this configuration is per MVPN, multiple VPNs can have identical policies configured, each for its own VPN context.

When a multicast stream is mapped to a multi-stream data MDT policy, the data will traverse the default MDT first. The data delay timer is used to switch the data from the default MDT to the multi-stream data MDT.

When the multi-stream data MDT is deleted, the traffic switches back to the default MDT.

In some cases when a new multi-stream data MDT is configured which is better suited, some streams might prefer this new multi-stream data MDT. To switch, the traffic will switch to the default MDT and then to the new multi-stream data MDT.

MDT data can be configured as SSM or ASM.

There can be multiple multi-stream policies assigned to a single data MDT configuration. In this case, the policy acts as a link list, the first (lowest index) matched multi-stream policy is used for that specific stream. The following are the rules of matching a multi-stream data MDT to (C-S, C-G) mapping on a source PE (in order):

1. The multi-stream policy is evaluated to enable the feature in service when per-(C-S, C-G) configuration is present, starting from the lowest numerical policy index (only entries that are not shutdown are evaluated). The first multi-stream data MDT (the lowest policy index) the (C-S, C-G) maps to is selected.
2. If (C-S, C-G) does not map to any of the multi-stream data MDTs, per-(C-S, C-G) single data MDTs are used if a (C-S, C-G) maps to an existing MDT based on data-thresholds.
3. The default MDT is used if no policy matches to a data MDT.
4. When a (C-S, C-G) arrives, it start on the default MDT but can switch to a data MDT when the data delay interval expires. It does not check the data-threshold before switching.

5. When going from one multi-stream data MDT to a more suitable one, the traffic will first switch to the default MDT and then switch to the new multi-stream data MDT based on the data-delay-interval.

When a data MDT tunnel fails, all (C-S, C-G)s using this multi-stream data MDT move to the default MDT, and the groups move back to the data MDT when it is restored.

### 3.4.7 MVPN (NG-MVPN) Upstream Multicast Hop Fast Failover

MVPN upstream PE or P node fast failover detection method is supported with RSVP P2MP I-PMSI only. A receiver PE achieves fast upstream failover based on the capability to subscribe multicast flow from multiple UMH nodes and the capability to monitor the health of the upstream PE and intermediate P nodes using an unidirectional multi-point BFD session running over the provider tunnel.

A receiver PE subscribes multicast flow from multiple upstream PE nodes to have active redundant multicast flow available during failure of primary flow. Active redundant multicast flow from standby upstream PE allows instant switchover of multicast flow during failure of primary multicast flow.

Faster detection of multicast flow failure is achieved by keeping track of unidirectional multi-point BFD sessions enabled on the provider tunnel. Multi-point BFD sessions must be configured with 10 ms transmit interval on sender (root) PE to achieve sub-50ms fast failover on receiver (leaf) PE.

UMH **tunnel-status** selection option must be enabled on the receiver PE for upstream fast failover. Primary and standby upstream PE pairs must be configured on the receiver PE to enable active redundant multicast flow to be received from the standby upstream PE.

## 3.4.8 Multicast VPN Extranet

Multicast VPN extranet distribution allows multicast traffic to flow across different routing instances. A routing instance that received a PIM/IGMP JOIN but cannot reach source of multicast source directly within its own instance is selected as receiver routing instance (receiver C-instance). A routing instance that has source of multicast stream and accepts PIM/IGMP JOIN from other routing instances is selected as source routing instance (source C-instance). A routing instance that does not have either source or receivers but is used in the core is selected as a transit instance (transit P-instance). The following subsections detail supported functionality.

### 3.4.8.1 Multicast Extranet for Rosen MVPN for PIM SSM

Multicast extranet is supported for Rosen MVPN with MDT SAFI. Extranet is supported for IPv4 multicast stream for default and data MDTs (PIM and IGMP).

The following extranet cases are supported:

- local replication into a receiver VRF from a source VRF on a source PE
- Transit replication from a source VRF onto a tunnel of a transit core VRF on a source PE. A source VRF can replicate its streams into multiple core VRFs as long as any given stream from source VRF uses a single core VRF (the first tunnel in any core VRF on which a join for the stream arrives). Streams with overlapping group addresses (same group address, different source address) are supported in the same core VRF.
- remote replication from source/transit VRF into one or more receiver VRFs on receiver PEs
- multiple replications from multiple source/transit VRFs into a receiver VRF on receiver PEs

Rosen MVPN extranet requires routing information exchange between the source VRF and the receiver VRF based on route export/import policies:


- Routing information for multicast sources must be exported using an RT export policy from the source VRF instance.
- Routing information must be imported into the receiver/transit VRF instance using an RT import policy.

Caveats:

- The source VRF instance and receiver VRF instance of extranet must exist on a common PE node (to allow local extranet mapping).
- SSM translation is required for IGMP (C-\*, C-G).
- An I-PMSI route cannot be imported into multiple VPRNs, and NG-MVPN routes do not need to be imported into multiple VPRNs.

In [Figure 30](#), VPRN-1 is the source VPRN instance and VPRN-2 and VPRN-3 are receiver VPRN instances. The PIM/IGMP JOIN received on VPRN-2 or VPRN-3 is for (S1, G1) multicast flow. Source S1 belongs to VPRN-1. Due to the route export policy in VPRN-1 and the import policy in VPRN-2 and VPRN-3, the receiver host in VPRN-2 or VPRN-3 can subscribe to the stream (S1, G1).

**Figure 30 Multicast VPN Traffic Flow**



al\_0164

### 3.4.8.2 Multicast Extranet for NG-MVPN for PIM SSM

Multicast extranet is supported for ng-MVPN with IPv4 RSVP-TE and mLDP I-PMSIs and S-PMSIs including (C-\*, C-\*) S-PMSI support where applicable. Extranet is supported for IPv4 C-multicast traffic (PIM/IGMP joins).

The following extranet cases are supported:

- local replication into a receiver C-instance MVPN(s) on a source PE from a source P-instance MVPN
- remote replication from P-instance MVPN into one or more receiver C-instance MVPNs on receiver PEs
- multiple replications from multiple source/transit P-instance MVPNs into a receiver C-instance MVPN on receiver PEs
- Transit replication on Source PE is not supported.

Multicast extranet for ng-MVPN, similarly to extranet for Rosen MVPN, requires routing information exchange between source ng-MVPN and receiver ng-MVPN based on route export and import policies. Routing information for multicast sources is exported using an RT export policy from a source ng-MVPN instance and imported into a receiver ng-MVPN instance using an RT import policy. S-PMSI/I-PMSI establishment and C-multicast route exchange occurs in a source ng-MVPN P-instance only (import and export policies are not used for MVPN route exchange). Sender-only functionality must not be enabled for the source/transit ng-MVPN on the receiver PE. It is recommended to enable receiver-only functionality on a receiver ng-MVPN instance.

Caveats:


- Source P-instance MVPN and receiver C-instance MVPN must reside on the receiver PE (to allow local extranet mapping).
- SSM translation is required for IGMP (C-\*, C-G).

### 3.4.8.3 Multicast Extranet with Per-group Mapping for PIM SSM

In some deployments, such as IPTV or wholesale multicast services, it may be desirable to create one or more transit MVPNs to optimize delivery of multicast streams in the provider core. [Figure 31](#) represents a sample deployment model.



**Figure 31 Source PE Transit Replication and Receiver PE Per-group SSM Extranet Mapping**



The architecture displayed in [Figure 31](#) requires a source routing instance MVPN to place its multicast streams into one or more transit core routing instance MVPNs (each stream mapping to a single transit core instance only). It also requires receivers within each receiver routing instance MVPN to know which transit core routing instance MVPN they need to join for each of the multicast streams. To achieve this functionality, transit replication from a source routing instance MVPN onto a tunnel of a transit core routing instance MVPN on a source PE (see earlier sub-sections for MVPN topologies supporting transit replication on source PEs) and per-group mapping of multicast groups from receiver routing instance MVPNs to transit core routing instance MVPNs (as defined below) are required.

For per-group mapping on a receiver PE, the operator must configure a receiver routing instance MVPN per-group mapping to one or more source/transit core routing instance MVPNs. The mapping allows propagation of PIM joins received in the receiver routing instance MVPN into the core routing MVPN instance defined by the map. All multicast streams sourced from a single multicast source address are always mapped to a single core routing instance MVPN for a given receiver routing instance MVPN (multiple receiver MVPNs can use different core MVPNs for groups from the same multicast source address). If per-group map in receiver MVPN maps multicast streams sourced from the same multicast source address to multiple core routing instance MVPNs, then the first PIM join processed for those streams selects the core routing instance MVPN to be used for all multicast streams from a given source address for this receiver MVPN. PIM joins for streams sourced from the

source address not carried by the selected core VRF MVPN instance will remain unresolved. When a PIM join/prune is received in a receiver routing instance MVPN with per-group mapping configured, if no mapping is defined for the PIM join's group address, non-extranet processing applies when resolving how to forward the PIM join/prune.

The main attributes for per-group SSM extranet mapping on receiver PE include support for:

- Rosen MVPN with MDT SAFI. \* RFC6513/6514 NG-MVPN with IPv4 RSVP-TE/mLDP in P-instance (a P-instance tunnel must be in the same VPRN service as multicast source)
- IPv4 PIM SSM
- IGMP (C-S, C-G), and for IGMP (C-\*, C-G) using SSM translation
- a receiver routing instance MVPN to map groups to multiple core routing instance MVPNs
- in-service changes of the map to a different transit/source core routing instance (this is service affecting)

Caveats:


- When a receiver routing instance MVPN is on the same PE as a source routing instance MVPN, basic extranet functionality and not per-group (C-S, C-G) mapping must be configured (extranet from receiver routing instance to core routing instance to source routing instance on a single PE is not supported).
- Local receivers in the core routing instance MVPN are not supported when per-group mapping is deployed.
- Receiver routing instance MVPN that has per-group mapping enabled cannot have tunnels in its OIF lists.
- Per-group mapping is blocked if GRT/VRF extranet is configured.

#### **3.4.8.4 Multicast GRT-source/VRF-receiver Extranet with Per Group Mapping for PIM SSM**

Multicast GRT-source/VRF-receiver (GRT/VRF) extranet with per-group mapping allows multicast traffic to flow from GRT into VRF MVPN instances. A VRF routing instance that received a PIM/IGMP join but cannot reach the source multicast stream directly within its own instance is selected as receiver routing instance. A GRT instance that has sources of multicast streams and accepts PIM joins from other VRF MVPN instances is selected as source routing instance.

Figure 32 depicts a sample deployment.

**Figure 32 GRT/VRF Extranet**



Routing information is exchanged between GRT and VRF receiver MVPN instances of extranet by enabling grt-extranet under a receiver MVPN PIM configuration for all or a subset of multicast groups. When enabled, multicast receivers in a receiver routing instance(s) can subscribe to streams from any multicast source node reachable in GRT source instance.


The main feature attributes are:

- GRT/VRF extranet can be performed on all streams or on a configured group of prefixes within a receiver routing instance.
- GRT instance requires Classic Rosen multicast.
- IPv4 PIM joins are supported in receiver VRF instances.
- Local receivers using IGMP: (C-S, C-G) and (C-\*, C-G) using SSM translation are supported.
- The feature is blocked if a per-group mapping extranet is configured in receiver VRF.

### 3.4.8.5 Multicast Extranet with Per-group Mapping for PIM ASM

Multicast extranet with per-group mapping for PIM ASM allows multicast traffic to flow from a source routing instance to a receiver routing instance when a PIM ASM join is received in the receiver routing instance.

Figure 33 depicts PIM ASM extranet map support.

**Figure 33 Multicast Extranet with per group PIM ASM mapping**

PIM ASM extranet is achieved by local mapping from the receiver to source routing instances on a receiver PE. The mapping allows propagation of anycast RP PIM register messages between the source and receiver routing instances over an auto-created internal extranet interface. This PIM register propagation allows the receiver routing instance to resolve PIM ASM joins to multicast source(s) and to propagate PIM SSM joins over an auto-created extranet interface to the source routing instance. PIM SSM joins are then propagated towards the multicast source within the source routing instance.

The following MVPN topologies are supported:

- Rosen MVPN with MDT SAFI: a local replication on a source PE and multiple-source/multiple-receiver replication on a receiver PE
- RFC 6513/6514 NG-MVPN (including RFC 6625 (C-\*, C-\*) wildcard S-PMSI): a local replication on a source PE and a multiple source/multiple receiver replication on a receiver PE
- Extranet for GRT-source/VRF receiver with a local replication on a source PE and a multiple-receiver replication on a receiver PE
- Locally attached receivers are supported without SSM mapping.

To achieve the extranet replication, the operator must configure:

- local PIM ASM mapping on a receiver PE from a receiver routing instance to a source routing instance (**config>service>vrpn>mvpn>rpf-select>core-mvpn** or **config>service>vrpn>pim>grt-extranet** as applicable)

- an anycast RP mesh between source and receiver PEs in the source routing instance


Caveats:

- This feature is supported for IPv4 multicast.
- The multicast source must reside in the source routing instance the ASM map points to on a receiver PE (the deployment of transit replication extranet from source instance to core instance on Source PE with ASM map extranet from receiver instance to core instance on a receiver PE is not supported).
- A given multicast group can be mapped in a receiver routing instance using either PIM SSM mapping or PIM ASM mapping but not both.
- A given multicast group cannot map to multiple source routing instances.

### **3.4.9 Non-Congruent Unicast and Multicast Topologies for Multicast VPN**

Operators that prefer to keep unicast and multicast traffic on separate links in a network have the option to maintain two separate instances of the route table (unicast and multicast) per VPRN.

Multicast BGP can be used to advertise separate multicast routes using Multicast NLRI (SAFI 2) on PE-CE link within VPRN instance. Multicast routes maintained per VPRN instance can be propagated between PE-PE using BGP Multicast-VPN NLRI (SAFI 129).

**Figure 34** Incongruent Multicast and Unicast Topology for Non-Overlapping Traffic Links

al\_0074

SR OS supports option to perform RPF check per VPRN instance using multicast route table, unicast route table or both.

Non-congruent unicast and multicast topology is supported with NG-MVPN. Draft Rosen is not supported.

### 3.4.10 Automatic Discovery of Group-to-RP Mappings (Auto-RP)


Auto-RP is a proprietary group discovery and mapping mechanism for IPv4 PIM that is described in *cisco-ipmulticast/pim-autorp-spec*, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast*. The functionality is similar to the IETF standard bootstrap router (BSR) mechanism that is described in RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*, to dynamically learn about the availability of Rendezvous Points (RPs) in a network. When a router is configured as an RP-mapping agent with the **pim>rp>auto-rp-discovery** command, it listens to the CISCO-RP-ANNOUNCE (224.0.1.39) group and caches the announced mappings. The RP-mapping agent then periodically sends out RP-mapping packets to the CISCO-RP-DISCOVERY (224.0.1.40) group. SR OS supports version 1 of the auto-RP specification, so the ability to deny RP-mappings by advertising negative group prefixes is not supported.

PIM dense-mode (PIM-DM) as described in RFC 3973, *Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)*, is used for the auto-RP groups to support multihoming and redundancy. The RP-mapping agent supports announcing, mapping, and discovery functions; candidate RP functionality is not supported.

Auto-RP is supported for IPv4 in multicast VPNs and in the global routing instance. Either BSR or auto-RP for IPv4 can be configured; the two mechanisms cannot be enabled together. BSR for IPv6 and auto-RP for IPv4 can be enabled together. In a multicast VPN, auto-RP cannot be enabled together with sender-only or receiver-only multicast distribution trees (MDTs), or wildcard S-PMSI configurations that could block flooding.

### 3.4.11 IPv6 MVPN Support

IPv6 multicast support in SR OS allows operators to offer customers IPv6 multicast MVPN service. An operator utilizes IPv4 mLDP or RSVP-TE core to carry IPv6 c-multicast traffic inside IPv4 mLDP or RSVP-TE provider tunnels (p-tunnels). The IPv6 customer multicast on a given MVPN can be blocked, enabled on its own or in addition to IPv4 multicast per PE or per interface. When both IPv4 and IPv6 multicast is enabled for a given MVPN, a single tree is used to carry both IPv6 and IPv4 traffic. [Figure 35](#) shows an example of an operator with IPv4 MPLS backbone providing IPv6 MVPN service to Customer 1 and Customer 2.

**Figure 35 IPv6 MVPN Example**

al\_0168

SR OS IPv6 MVPN multicast implementation provides the following functionality:

- IPv6 C-PIM-SM (ASM and SSM)
- MLDv1 and MLDv2
- SSM mapping for MLDv1
- I-PMSI and S-PMSI using IPv4 P2MP mLDP p-tunnels
- I-PMSI and S-PMSI using IPv4 P2MP RSVP p-tunnels
- BGP auto-discovery
- PE-PE transmission of C-multicast routing using BGP mvpn-ipv6 address family
- IPv6 BSR/RP functions on functional par with IPv4 (auto-RP using IPv4 only)
- Embedded RP
- Inter-AS Option A




The following known caveats exist for IPv6 MVPN support:

- Non-congruent topologies are not supported
- IPv6 is not supported in MCAC
- If IPv4 and IPv6 multicast is enabled, per-MVPN multicast limits apply to entire IPv4 and IPv6 multicast traffic as it is carried in a single PMSI. For example IPv4 AND IPv6 S-PMSIs are counted against a single S-PMSI maximum per MVPN.
- IPv6 Auto-RP is not supported

### 3.4.12 Multicast Core Diversity for Rosen MDT\_SAFI MVPNs

Figure 36 depicts Rosen MVPN core diversity deployment:

**Figure 36 Multicast Core Diversity**



Core diversity allows operator to optionally deploy multicast MVPN in either default IGP instance or one of two non-default IGP instances to provide, for example, topology isolation or different level of services. The following describes main feature attributes:

1. Rosen MVPN IPv4 multicast with MDT SAFI is supported with default and data MDTs.
2. Rosen MVPN can use a non-default OSPF or ISIS instance (using their loopback addresses instead of a system address).

3. Up to 3 distinct core instances are supported: system + 2 non-default OSPF instances – referred as “red” and “blue” below.
4. The BGP Connector also uses non-default OSPF loopback as NH, allowing Inter-AS Option B/C functionality to work with Core diversity as well.
5. The feature is supported with CSC-VRPN.

On source PEs (PE1: UMH, PE2: UMH in [Figure 36](#)), an MVPN is assigned to a non-default IGP core instance as follows:

1. MVPN is statically pointed to use one of the non-default “red”/“blue” IGP instances loopback addresses as source address instead of system loopback IP.
2. MVPN export policy is used to change unicast route next-hop VPN address (no longer required as of SR OS Release 12.0.R4 - BGP Connector support for non-default instances).

The above configuration ensures that MDT SAFI and IP-VPN routes for the non-default core instance use non-default IGP loopback instead of system IP. This ensures PIM advertisement/joins run in the proper core instance and GRE tunnels for multicast can be set-up using and terminated on non-system IP.

If BGP export policy is used to change unicast route next-hop VPN address, unicast traffic must be forwarded in non-default “red” or “blue” core instance LDP or RSVP (terminating on non-system IP) must be used. GRE unicast traffic termination on non-system IP is not supported, and any GRE traffic arriving at the PE in “blue”, “red” instances destined to non-default IGP loopback IP will be forwarded to CPM (ACL or CPM filters can be used to prevent the traffic from reaching the CPM). This limitation does not apply if BGP connector attribute is used to resolve the multicast route.

No configuration is required on non-source PEs.


Feature caveats:

- VPRN instance must be shutdown to change the mdt-safi source-address. The CLI rollback that includes change of the auto-discovery is thus service impacting.
- To reset mdt-safi source-address to system IP, operator must first execute no auto-discovery (or auto-discovery default) then auto-discovery mdt-safi
- Configuring system IP as a source-address will consume one of the 2 IP addresses allowed, thus it should not be done.
- Operators must configure proper IGP instance loopback IP addresses within Rosen MVPN context and must configure proper BGP policies (prior to release 12.0.R4) for the feature to operate as expected. There is no verification that the address entered for MVPN provider tunnel source-address is such an address or is not a system IP address.

### 3.4.13 NG-MVPN Core Diversity

See [Figure 37](#) for an operational example of logical networks using Multi-Instance IGP.

**Figure 37** Logical Networks Using Multi-Instance IGP



sw0113

SR OS is being positioned in multi-instance IGP as a virtualization or migration strategy in numerous cases. One specific application is as a virtual LSR core whereby various topologies are created using separate IGP instances. Specifically, the migration to SR solutions requires the deployment of multi-instance IGP with service migrations. The objective is to more cleanly segregate protocols and service bindings to specific routing instances. NG-MVPN does not currently allow non-system loopbacks to be used for PMSI (for example, MVPN address family).


The ability to support binding MVPN with mLDP tunnels to different loopback interfaces is one of the main requirements. In addition, assigning these loopback interfaces to different IGP instances will create parallel NG-MVPN services, each running over a separate IGP instance.

This feature has two main components to it:

1. The ability of advertising a MVPN route, via a loopback interface, and generating an mLDP FEC with that loopback interface.
2. The ability of creating multiple IGP instances and assigning a loopback to each IGP instance. This, combined with the component above, will create parallel NG-MVPN services on different IGP instances.

### 3.4.13.1 NG-MVPN to Loopback Interface

**Figure 38 NG-MVPN Setup Via Loopback Interface**



SystemIP is usually used for management purposes and, as such, it might be desired to create services to a separate loopback interface. Due to security concerns, many operators do not want to create services to the systemIP address.

The first portion of this feature will allow MVPN routes be advertised with a nexthop specific loopback.


This can be achieved via two methods:

1. Having iBGP session to the loopback interface of the peer, and using the corresponding local loopback for local-address.
2. Having iBGP session to the systemIP but using policies to change the nexthop of the AD route to the corresponding loopback interface.

After the AD routes are advertised via the loopback interface as nexthop, PIM will generate an mLDP FEC for the loopback interface.

The preceding configuration will allow an NG-MVPN to be established via a specific loopback.


**Figure 39 Intra-as Basic Opaque FEC to Loopback Interface**



It should be noted that this setup will work in a single area or multi-area through an ABR.

### 3.4.13.2 NG-MVPN Core Diversity

**Figure 40 Core Diversity with Parallel NG-MVPN Services on Parallel IGP Instances**



In [Figure 40](#), Red and Blue networks correspond to separate IGP instances tied to separate loopback interfaces. In core diversity, MVPN services on the LER are bound to a domain by advertising MP-BGP routes, with next hop set to the appropriate loopback address, and having the receiving LERs resolve those BGP next hops, based on the corresponding IGP instance. All subsequent MVPN BGP routing exchanges must set and resolve the BGP next hop with the configured non-system loopback.

With this feature, an MP-BGP next hop would resolve to a link LDP label which is indirectly associated with a given IGP instance. Services which advertise BGP routes with next hop *red* loopback would result in traffic flowing over the *Red* network using LDP and *blue* loopback would result in traffic flowing over *Blue* network. See [Figure 40](#).

Most importantly, the common routes in each IGP instance must have a unique and active label. This is required to ensure that the same route advertised in two different IGP domains do not resolve to the same label. LDP *local-lsr-id* can be used to ensure FECs and label mapping be advertised via the right instance of IGP.

Core diversity allows an operator to optionally deploy multicast NG-MVPN in either default IGP instance or one of the non-default IGP instances to provide, for example, topology isolation or different level of services. The following describes the main feature attributes:

- NG- MVPN can use IPv4/IPv6 multicast.
- NG-MVPN can use a non-default OSPF or ISIS instance.



**Note:** This is accomplished by using their loopback addresses instead of a system address.

- The BGP Connector also uses non-default OSPF loopback as NH, via two methods:
  - a. Setting the BGP local-address to a loopback interface IP address
  - b. Creating a routing policy to set the NH of a MVPN AD route to the corresponding loopback IP
- RSVP-TE/LDP transport tunnels also use non-systemIP loopback for session creation. As an example, RSVP-TE p2mp LSPs would be created to non systemIP loopbacks and mLDP will create a session via *local-lsr-id*.
- Need to add the source-address for mvpn auto-discover default.

On the source PEs, a NG-MVPN is assigned to a non-default IGP core instance as follows:


1. NG-MVPN is statically pointed to use one of the non-default *red/blue* IGP instances loopback addresses as source address instead of system loopback IP.
2. MVPN export policy is used to change unicast route next-hop VPN address (BGP Connector support for non-default instances).
3. Alternatively, the BGP local-address can be set to the correct loopback interface assigned for that specific instance.

The preceding configuration ensures that MVPN-IPv4/v6 and IP-VPN routes for the non-default core instance use non-default IGP loopback instead of system IP. This ensures MVPN-IPv4/v6 advertisement/joins run in the proper core instance and mLDP and P2MP RSVP tunnels (I-PMSI and S-PMSI) for multicast can be set-up using and terminating on non-system IP.

If BGP export policy is used to change unicast route next-hop VPN address, than unicast traffic must be forwarded in non-default *red* or *blue* core instance LDP or RSVP (terminating on non-system IP) must be used.

### 3.4.13.3 Configuration Example

**Figure 41** Core Diversity with Parallel NG-MVPN Services on Parallel IGP Instances



sw0117

In Figure 41, there are three IGP instances, the default Instance 0, Instance 1, and Instance 2. Each instance will bind to its own loopback interface. For Instance 0, it will be the systemIP *system Interface* used as loopback. LDP, RSVP and MP-BGP need to run between the corresponding loopback associated with each instance.

For example, for the blue Instance 2, both MP-BGP and LDP need to be configured to its corresponding loopback *loopback2* and the next-hop for BGP MVPN-IP4/IPv6 and the VPN-IP4/IPv6 need to be *loopback2*.

From a configuration point of view, the following steps need to be taken:

1. For MLDP, configure LDP with *local-lsr-id* with loopback interface of instance 2 *loopback2*:

```
*A:SwSim14>config>router>ldp# info

interface-parameters
 interface "2ROOT" dual-stack
 ipv4
 local-lsr-id interface-name "loopback2"
 no shutdown
 exit
 no shutdown
exit
```

2. Enable the source-address for default auto discover as follows:

```
*A:Dut-A>config>service>vprn 2
 mvpn auto-discovery default source-address LOOPBack1
 vrf-export "vprnexpl00"
*A:Dut-A>config>service>vprn 3
 mvpn auto-discovery default source-address LOOPBack2
 vrf-export "vprnexpl01"
```

3. Define community *vprnXXXX* for each VPRN using non-default core-instance and define a policy to tag each VPRN with either a *blue* or *red* standard community attribute:

```
*A:Dut-A>config>router>policy-options# info
 community "vprn2" members "target:70:70"
 community "vprn3" members "target:80:80"
 policy-statement "vprnexp2"
 entry 10
 from
 protocol direct
 exit
 action accept
 community add "vprn2" "red"
 exit
 exit
 policy-statement "vprnexp3"
 entry 10
 from
 protocol direct
 exit
 action accept
 community add "vprn3" "blue"
 exit
 exit
 exit
```

4. Define a single global BGP policy to change next hop for red/blue MVPNs:

```
*A:Dut-A>config>router>policy-options# info
 policy-statement "MVPN_CoreDiversity_Exp"
 entry 10
 from
```



```

 community "red"
 exit
 to
 protocol bgp-vpn
 exit
 action accept
 next-hop loopback1
 exit
exit
entry 20
 from
 community "blue"
 exit
 to
 protocol bgp-vpn
 exit
 action accept
 next-hop loopback2
 exit
exit
exit

```

5. Configure BGP default MVPN export in the group as required:

```
configure router bgp group "mvpn" export "MVPN_CoreDiveristy_Exp"
```

6. Configure each VPRN to use proper IGP source address and proper VRF export policy:

```


*A:Dut-A>config>service>vprn 2
 mvpn auto-discovery default source-address loopback1
 vrf-export "vprnexp2"
*A:Dut-A>config>service>vprn 3
 mvpn auto-discovery default source-address loopback2
 vrf-export "vprnexp3"

```

### 3.4.14 NG-MVPN Multicast Source Geo-Redundancy

Multicast source geo-redundancy is targeted primarily for MVPN deployments for multicast delivery services like IPTV. The solution allows operators to configure a list of geographically dispersed redundant multicast sources (with different source IPs) and then, using configured BGP policies, ensure that each Receiver PE (a PE with receivers in its C-instance) selects only a single, most-preferred multicast source for a given group from the list. Although the data may still be replicated in P-instance (each multicast source sends (C-S, C-G) traffic onto its I-PMSI tree or S-PMSI tree), each Receiver PE only forwards data to its receivers from the preferred multicast source. This allows operators to support multicast source geo-redundancy without the replication of traffic for each (C-S, C-G) in the C-instance while allowing fast recovery of service when an active multicast source fails.

Figure 42 shows an operational example of multicast source geo-redundancy:

**Figure 42 Preferred Source Selection for Multicast Source Geo-Redundancy**

Operators can configure a list of prefixes for multicast source redundancy per MVPN on Receiver PEs:

- Up to 8 multicast source prefixes per VPRN are supported.
- Any multicast source that is not part of the source prefix list is treated as a unique source and automatically joined in addition to joining the most preferred source from the redundant multicast source list.

A Receiver PE selects a single, most-preferred multicast source from the list of pre-configured sources for a given MVPN during (C-\*, C-G) processing as follows:

- A single join for the group is sent to the most preferred multicast source from the operator-configured multicast source list. Joins to other multicast sources for a given group are suppressed. Operator can see active and suppressed joins on a Receiver PE. Although a join is sent to a single multicast source only, (C-S, C-G) state is created for every source advertising Type-5 S-A route on the Receiver PE.
- The most preferred multicast source is a reachable source with the highest local preference for Type-5 SA route based on the BGP policy, as described later in this section.

- On a failure of the preferred multicast source or when a new multicast source with a better local preference is discovered, Receiver PE will join the new most-preferred multicast source. The outage experienced will depend on how quickly Receiver PE receives Type-5 S-A route withdrawal or loses unicast route to multicast source, and how quickly the network can process joins to the newly selected preferred multicast source(s).
- Local multicast sources on a Receiver PE are not subject to the most-preferred source selection, regardless of whether they are part of redundant source list or not.

BGP policy on Type-5 SA advertisements is used to determine the most preferred multicast source based on the best local preference as following:

- Each Source PE (a PE with multicast sources in its C-instance) tags Type-5 SA routes with a unique standard community attribute using global BGP policy or MVPN vrf-export policy. Depending on multicast topology, the policy may require source-aware tagging in the policy. Either all MVPN routes or Type 5 SA routes only can be tagged in the policy (new attribute **mvpn-type 5**).
- Each receiver PE has a BGP VRF import policy that sets local preference using match on Type-5 SA routes (new attribute **mvpn-type 5**) and standard community attribute value (as tagged by the Source PEs). Using policy statements that also include group address match, allows receiver PEs to select the best multicast source per group. The BGP VRF import policy must be applied as **vrf-import** under *config>service>vprn>mvpn* context. It must have default-action **accept** specified, or all MVPN routes other than those matched by specified entries will be rejected. In addition, it must have **vrf-target** as a community match condition, because **vrf-target mvpn** configuration is ignored when **vrf-import** policy is defined.

Operators can change redundant source list or BGP policy affecting source selection in service. If such a change of the list/policy results in a new preferred multicast source election, make-before-break is used to join the new source and prune the previously best source.

For the proper operations, MVPN multicast source geo-redundancy requires the router:


- To maintain the list of eligible multicast sources on Receiver PEs, Source PE routers must generate Type-5 S-A route even if the Source PE sees no active joins from any receiver for a given group.
- To trigger a switch from a currently active multicast source on a Receiver PE, Source PE routers must withdraw Type-5 S-A route when the multicast source fails or alternatively unicast route to multicast source must be withdrawn or go down on a Receiver PE.

MVPN multicast source redundancy solutions is supported for the following configurations only. Enabling the feature in unsupported configuration must be avoided:

1. NG-MVPN with RSVP-TE or mLDP or PIM with BGP c-multicast signaling in P-instance. Both I-PMSI and S-PMSI trees are supported.
2. IPv4 and IPv6 (C-\*, C-G) PIM ASM joins in the C-instance.
3. Both **intersite-shared enabled** and **disabled** are supported. For **intersite-shared enabled**, operators must enable generation of Type-5 S-A routes even in the absence of receivers seen on Source PEs (intersite-shared persistent-type5-adv must be enabled).
4. The Source PEs must be configured as a sender-receiver, the Receiver PEs can be configured as a sender-receiver or a receiver-only.
5. The RP(s) must be on the Source PE(s) side. Static RP, anycast-RP, embedded-RP types are supported.
6. UMH redundancy can be deployed to protect Source PE to any multicast source. When deployed, UMH selection is executed independently of source selection after the most preferred multicast source had been chosen. Supported **umh-selection** options include: **highest-ip**, **hash-based**, **tunnel-status** (not supported for IPv6), and **unicast-rt-pref**.

### 3.4.15 Multicast Core Diversity for Rosen MDT SAFI MVPNs

[Figure 43](#) shows a Rosen MVPN core diversity deployment.



- MVPN export policy is used to change unicast route next-hop VPN address.
- BGP Connector support for non-default instances.

The configuration shown above ensures that MDT SAFI and IP-VPN routes for the non-default core instance use non-default IGP loopback instead of system IP. This ensures PIM advertisement/joins run in the proper core instance and GRE tunnels for multicast can be set-up using and terminated on non-system IP. If BGP export policy is used to change unicast route next-hop VPN address instead of BGP Connector attribute-based processing and unicast traffic must be forwarded in non-default core instances 1 or 2, LDP or RSVP (terminating on non-system IP) must be used. GRE unicast traffic termination on non-system IP is not supported and any GRE traffic arriving at the PE in instances 1 or 2, destined to non-default IGP loopback IP will be forwarded to CPM (ACL or CPM filters can be used to prevent the traffic from reaching the CPM).

No configuration is required on non-source PEs.

Known feature caveats include:

- VPRN instance must be shutdown to change the mdt-safi source-address. The CLI rollback that includes change of the auto-discovery is thus service impacting.
- To reset mdt-safi source-address to system IP, operator must first execute no auto-discovery (or auto-discovery default) then auto-discovery mdt-safi
- Configuring system IP as a source-address will consume one of the 2 IP addresses allowed, thus it should not be done.
- Operators must configure proper IGP instance loopback IP addresses within Rosen MVPN context and must configure proper BGP policies (prior to release 12.0R4) for the feature to operate as expected. There is no verification that the address entered for MVPN provider tunnel source-address is such an address or is not a system IP address.

### 3.4.16 Inter-AS MVPN

The Inter-AS MVPN feature allows set-up of Multicast Distribution Trees (MDTs) that span multiple Autonomous Systems (ASes). This section focuses on multicast aspects of the Inter-AS MVPN solution.

To support Inter-AS option for MVPNs, a mechanism is required that allows setup of Inter-AS multicast tree across multiple ASes. Due to limited routing information across AS domains, it is not possible to setup the tree directly to the source PE. Inter-AS VPN Option A does not require anything specific to inter-AS support as customer instances terminate on ASBR and each customer instance is handed over to the other AS domain via a unique instance. This approach allows operators to provide full isolation of ASes, but the solution is the least scalable case, as customer instances across the network have to exist on ASBR.

Inter-AS MVPN Option B allows operators to improve upon the Option A scalability while still maintaining AS isolation, while Inter-AS MVPN Option C further improves Inter-AS scale solution but requires exchange of Inter-AS routing information and thus is typically deployed when a common management exists across all ASes involved in the Inter-AS MVPN. The following sub-sections provide further details on Inter-AS Option B and Option C functionality.

### **3.4.16.1 BGP Connector Attribute**

BGP connector attribute is a transitive attribute (unchanged by intermediate BGP speaker node) that is carried with VPNv4 advertisements. It specifies the address of source PE node that originated the VPNv4 advertisement.

With Inter-AS MVPN Option B, BGP next-hop is modified by local and remote ASBR during re-advertisement of VPNv4 routes. On BGP next-hop change, information regarding the originator of prefix is lost as the advertisement reaches the receiver PE node.

BGP connector attribute allows source PE address information to be available to receiver PE, so that a receiver PE is able to associate VPNv4 advertisement to the corresponding source PE.

### **3.4.16.2 PIM RPF Vector**

In case of Inter-AS MVPN Option B, routing information towards the source PE is not available in a remote AS domain, since IGP routes are not exchanged between ASes. Routers in an AS other than that of a source PE, have no routes available to reach the source PE and thus PIM JOINS would never be sent upstream. To enable setup of MDT towards a source PE, BGP next-hop (ASBR) information from that PE's

MDT-SAFI advertisement is used to fake a route to the PE. If the BGP next-hop is a PIM neighbor, the PIM JOINS would be sent upstream. Otherwise, the PIM JOINS would be sent to the immediate IGP next-hop (P) to reach the BGP next-hop. Since the IGP next-hop does not have a route to source PE, the PIM JOIN would not be propagated forward unless it carried extra information contained in RPF Vector.

In case of Inter-AS MVPN Option C, unicast routing information towards the source PE is available in a remote AS PEs/ASBRs as BGP 3107 tunnels, but unavailable at remote P routers. If the tunneled next-hop (ASBR) is a PIM neighbor, the PIM JOINS would be sent upstream. Otherwise, the PIM JOINS would be sent to the immediate IGP next-hop (P) to reach the tunneled next-hop. Since the IGP next-hop does not have a route to source PE, the PIM JOIN would not be propagated forward unless it carried extra information contained in RPF Vector.

To enable setup of MDT towards a source PE, PIM JOIN thus carries BGP next hop information in addition to source PE IP address and RD for this MVPN. For option-B, both these pieces of information are derived from MDT-SAFI advertisement from the source PE. For option-C, both these pieces of information are obtained from the BGP tunneled route.

The RPF vector is added to a PIM join at a PE router when configure router **pim rpfv** option is enabled. P routers and ASBR routers must also have the option enabled to allow RPF Vector processing. If the option is not enabled, the RPF Vector is dropped and the PIM JOIN is processed as if the PIM Vector were not present.


For further details about RPF Vector processing please refer to [RFCs 5496, 5384 and 6513]

### 3.4.16.3 Inter-AS MVPN Option B

Inter-AS Option B is supported for Rosen MVPN PIM SSM using BGP MDT SAFI, PIM RPF Vector and BGP Connector attribute. The [Figure 44](#) depict set-up of a default MDT:



**Figure 44 Inter-AS Option B Default MDT Setup**




SR OS inter-AS Option B is designed to be standard compliant based on the following RFCs:

- RFC 5384 - The Protocol Independent Multicast (PIM) Join Attribute Format
- RFC 5496 - The Reverse Path Forwarding (RPF) Vector TLV
- RFC 6513 - Multicast in MPLS/BGP IP VPNs

The SR OS implementation was designed also to interoperate with older routers Inter-AS implementations that do not comply with the RFC 5384 and RFC 5496.

### 3.4.16.4 Inter-AS MVPN Option C

Inter-AS Option C is supported for Rosen MVPN PIM SSM using BGP MDT SAFI and PIM RPF Vector. [Figure 45](#) depicts a default MDT setup:

**Figure 45 Inter-AS Option C Default MDT Setup**

Additional caveats for Inter-AS MVPN Option B and C support are the following:

1. Inter-AS MVPN Option B is not supported with duplicate PE addresses.
2. For Inter-AS Option C, BGP 3107 routes are installed into unicast rtm (rtable-u), unless routes are installed by some other means into multicast rtm (rtable-m), and Option C will not build core MDTs, therefore, rpf-table is configured to rtable-u or both.
3. Additional Cisco interoperability notes are the following:
  - RFC 5384 - the Protocol Independent Multicast (PIM) Join Attribute Format
  - RFC 5496 - the Reverse Path Forwarding (RPF) Vector TLV
  - RFC 6513 - multicast in MPLS/BGP IP VPNs

The SR OS implementation was designed to inter-operate with Cisco routers Inter-AS implementations that do not comply with the RFC5384 and RFC5496.

When **configure router pim rpfv mvpn** option is enabled, Cisco routers need to be configured to include RD in an RPF vector using the following command: **ip multicast vrf vrf-name rpf proxy rd vector** for interoperability. When Cisco routers are not configured to include RD in an RPF vector, operator should configure SR OS router (if supported) using **configure router pim rpfv core mvpn**: PIM joins received can be a mix of core and mvpn RPF vectors.

### 3.4.16.5 NG-MVPN Non-segmented Inter-AS Solution


This feature allows multicast services to use segmented protocols and span them over multiple autonomous systems (ASs), as done in unicast services. As IP VPN or GRT services span multiple IGP areas or multiple ASs, either due to a network designed to deal with scale or as result of commercial acquisitions, operators may require Inter-AS VPN (unicast) connectivity. For example, an Inter-AS VPN can break the IGP, MPLS and BGP protocols into access segments and core segments, allowing higher scaling of protocols by segmenting them into their own islands. SR OS also allows for similar provision of multicast services and for spanning these services over multiple IGP areas or multiple ASs.

For multicast VPN (MVPN), SR OS previously supported Inter-AS Model A/B/C for Rosen MVPN; however, when MPLS was used, only Model A was supported for Next Generation Multicast VPN (NG-MVPN) and d-mLDP signaling.


For unicast VPRNs, the Inter-AS or Intra-AS Option B and C breaks the IGP, BGP and MPLS protocols at ABR routers (in case of multiple IGP areas) and ASBR routers (in case of multiple ASs). At ABR and ASBR routers, a stitching mechanism of MPLS transport is required to allow transition from one segment to next, as shown in [Figure 46](#) and [Figure 47](#).

In [Figure 46](#), the Service Label (S) is stitched at the ASBR routers.

**Figure 46 Unicast VPN Option B with Segmented MPLS**



In [Figure 47](#), the 3107 BGP Label Route (LR) is stitched at ASBR1 and ASBR3. At ASBR1, the LR1 is stitched with LR2, and at ASBR3, the LR2 is stitched with TL2.

**Figure 47 Unicast VPN Option C with Segmented MPLS**

Previously, in case of NG-MVPN, segmenting an LDP MPLS tunnel at ASBRs or ABRs was not possible. As such, RFC 6512 and 6513 used a non-segmented mechanism to transport the multicast data over P-tunnels end-to-end through ABR and ASBR routers. The signaling of LDP needed to be present and possible between two ABR routers or two ASBR routers in different ASs.



**Note:** For unicast VPNs, it was usually preferred to only have eBGP between ASBR routers. The non-segmented behavior of d-mLDP would have broken this by requiring LDP signaling between ASBR routers.

SR OS now has d-mLDP non-segmented intra-AS and inter-AS signaling for NG-MVPN and GRT multicast. The non-segmented solution for d-mLDP is possible for inter-ASs as Option B and C.

### 3.4.16.5.1 Non-Segmented d-mLDP and Inter-AS VPN

There are three types of VPN Inter-AS solutions:

- [Inter-AS Option A](#)
- [Inter-AS Option B](#)
- [Inter-AS Option C](#)

Options B and C use recursive opaque types 8 and 7 respectively, from [Table 37](#).

**Table 37 Recursive Opaque Types**

| Opaque Type | Opaque Name                   | RFC      | SR OS Use                        |
|-------------|-------------------------------|----------|----------------------------------|
| 1           | Basic Type                    | RFC 6388 | VPRN Local AS                    |
| 3           | Transit IPv4                  | RFC 6826 | IPv4 multicast over mLDP in GRT  |
| 4           | Transit IPv6                  | RFC 6826 | IPv6 multicast over mLDP in GRT  |
| 7           | Recursive Opaque (Basic Type) | RFC 6512 | Inter-AS Option C MVPN over mLDP |
| 8           | Recursive Opaque (VPN Type)   | RFC 6512 | Inter-AS Option B MVPN over mLDP |

### Inter-AS Option A

In Inter-AS Option A, ASBRs communicate using VPN access interfaces, which need to be configured under PIM for the two ASBRs to exchange multicast information.

### Inter-AS Option B


The recursive opaque type used for Inter-AS Option B is the Recursive Opaque (VPN Type), shown as opaque type 8 in [Table 37](#).

### Inter-AS Option B Support for NG-MVPN

Inter-AS Option B requires additional processing on ASBR routers and recursive FEC encoding than that of Inter-AS Option A. Because BGP adjacency is not e2e, ASBRs must cache and use a PMSI route to build the tree. For that, mLDP recursive FEC must carry RD information—thus, VPN recursive FEC is required (opaque type 8).

In Inter-AS Option B, the PEs in two different ASs do not have their system IP address in the RTM. As such, for NG-MVPN, a recursive opaque value in mLDP FEC is required to signal the LSP to the first ASBR in the local AS path.

Because the system IPs of the peer PEs (Root-1 and Root-2) are not installed on the local PE (leaf), it is possible to have two PEs in different ASs with same system IP address, as shown in [Figure 48](#). However, SR OS does not support this topology. The system IP address of all nodes (root or leaf) in different ASs must be unique.

**Figure 48** Identical System IP on Multiple PEs (Option B)

For inter-AS Option B and NG-MVPN, SR OS as a leaf does not support multiple roots in multiple ASs with the same system IP and different RDs; however, the first root that is advertised to an SR OS leaf will be used by PIM to generate an MLDP tunnel to this actual root. Any dynamic behavior after this point, such as removal of the root and its replacement by a second root in a different AS, is not supported and the SR OS behavior is nondeterministic.


### I-PMSI and S-PMSI Establishment

I-PMSI and S-PMSI functionality follows RFC 6513 section 8.1.1 and RFC 6512 sections 3.1 and 3.2.1. For routing, the same rules as for GRT d-MLDP use case apply, but the VRR Route Import External community now encodes the VRF instance in the local administrator field.

Option B uses an outer opaque of type 8 and inter opaque of type 1 (see [Table 37](#)).

[Figure 49](#) depicts the processing required for I-PMSI and S-PMSI Inter-AS establishment.

**Figure 49 Non-segmented mLDP PMSI Establishment (Option B)**



1031

For non-segmented mLDP trees, A-D procedures follow those of the Intra-AS model, with the exception that NO EXPORT community must be excluded; LSP FEC includes mLDP VPN-recursive FEC.

For I-PMSI on Inter-AS Option B:

- A-D routes must be installed by ASBRs and next-hop information is changed as the routes are propagated, as shown in [Figure 49](#).
- PMSI A-D routes are used to provide inter-domain connectivity on remote ASBRs.

On a receipt of an Intra-AS PMSI A-D route, PE2 resolves PE1's address (next-hop in PMSI route) to a labeled BGP route with a next-hop of ASBR3, because PE1 (Root-1) is not known via IGP. Because ASBR3 is not the originator of the PMSI route, PE2 sources an mLDP VPN recursive FEC with a root node of ASBR3, and an opaque value containing the information advertised by Root-1 (PE-1) in the PMSI A-D route, shown below, and forwards the FEC to ASBR 3 using IGP.

**PE-2 LEAF FEC: (Root ASBR3, Opaque value {Root: ROOT-1, RD 60:60, Opaque Value: P2MPLSP-ID xx})**

When the mLDP VPN-recursive FEC arrives at ASBR3, it notes that it is the identified root node, and that the opaque value is a VPN-recursive opaque value. Because Root-1 PE1 is not known via IGP, ASBR3 resolves the root node of the VPN-Recursive FEC using PMSI A-D (I or S) matching the information in the VPN-recursive FEC (the originator being PE1 (Root-1), RD being 60:60, and P2MP LSP ID xx). This yields ASBR1 as next hop. ASBR3 creates a new mLDP FEC element with a root node of ASBR1, and an opaque value being the received recursive opaque value, as shown below. ASBR then forwards the FEC using IGP.

**ASBR-3 FEC: {Root ASBR 1, Opaque Value {Root: ROOT-1, RD 60:60, Opaque Value: P2MPLSP-ID xx}}**

When the mLDP FEC arrives at ASBR1, it notes that it is the root node and that the opaque value is a VPN-recursive opaque value. As PE1's ROOT-1 address is known to ASBR1 through the IGP, no further recursion is required. Regular processing begins, using received Opaque mLDP FEC information.

**ASBR-1 FEC: {Root: ROOT-1, Opaque Value: P2MP LSP-ID xx}**



**Note:** VPN-Recursive FEC carries P2MPLSP ID. The P2MPLSP ID is used in addition to PE RD and Root to select a route to the mLDP root using the correct I-PMSI or S-PMSI route.


The functionality as described above for I-PMSI applies also to S-PMSI and (C-\*, C-\*) S-PMSI.

### C-multicast Route Processing

C-multicast route processing functionality follows RFC 6513 section 8.1.2 (BGP used for route exchange). The processing is analogous to BGP Unicast VPN route exchange described in [Figure 46](#) and [Figure 47](#). [Figure 50](#) shows C-multicast route processing with non-segmented mLDP PMSI details.



**Figure 50 Non-segmented mLDP C-multicast Exchange (Option B)**



1032

## Inter-AS Option C

In Inter-AS Option C, the PEs in two different ASs have their system IP address in the RTM, but the intermediate nodes in the remote AS do not have the system IP of the PEs in their RTM. As such, for NG-MVPN, a recursive opaque value in mLDP FEC is need to signal the LSP to the first ASBR in the local AS path.

The recursive opaque type used for Inter-AS Option B is the Recursive Opaque (Basic Type), shown as opaque type 7 in [Table 37](#).

## Inter-AS Option C support for NG-MVPN

For Inter-AS Option C, on a leaf PE, a route exists to reach root PE's system IP and, as ASBRs can use BGP unicast routes, recursive FEC processing using BGP unicast routes and *not* VPN recursive FEC processing using PMSI routes is required.


## I-PMSI and S-PMSI Establishment

I-PMSI and S-PMSI functionality follows RFC 6513 section 8.1.1 and RFC 6512 Section 2. The same rules as per the GRT d-mLDP use case apply, but the VRR Route Import External community now encodes the VRF instance in the local administrator field.

Option C uses an outer opaque of type 7 and inter opaque of type 1.

Figure 51 shows the processing required for I-PMSI and S-PMSI Inter-AS establishment.

**Figure 51 Non-segmented mLDP PMSI Establishment (Option C)**



1029

For non-segmented mLDP trees, A-D procedures follow those of the Intra-AS model, with the exception that NO EXPORT Community must be excluded; LSP FEC includes mLDP recursive FEC (and not VPN recursive FEC).

For I-PMSI on Inter-AS Option C:

- A-D routes are not installed by ASBRs and next-hop information is not changed in MVPN A-D routes.
- BGP-labeled routes are used to provide inter-domain connectivity on remote ASBRs.

On a receipt of an Intra-AS I-PMSI A-D route, PE2 resolves PE1's address (N-H in PMSI route) to a labeled BGP route with a next-hop of ASBR3, because PE1 is not known via IGP. PE2 sources an mLDP FEC with a root node of ASBR3, and an opaque value, shown below, containing the information advertised by PE1 in the I-PMSI A-D route.

**PE-2 LEAF FEC: {Root = ASBR3, Opaque Value: {Root: ROOT-1, Opaque Value: P2MP-ID xx}}**

When the mLDP FEC arrives at ASBR3, it notes that it is the identified root node, and that the opaque value is a recursive opaque value. ASBR3 resolves the root node of the Recursive FEC (ROOT-1) to a labeled BGP route with the next-hop of ASBR1, because PE-1 is not known via IGP. ASBR3 creates a new mLDP FEC element with a root node of ASBR1, and an opaque value being the received recursive opaque value.

**ASBR3 FEC: {Root: ASBR1, Opaque Value: {Root: ROOT-1, Opaque Value: P2MP-ID xx}}**


When the mLDP FEC arrives at ASBR1, it notes that it is the root node and that the opaque value is a recursive opaque value. As PE-1's address is known to ASBR1 through the IGP, no further recursion is required. Regular processing begins, using the received Opaque mLDP FEC information.

The functionality as described above for I-PMSI applies to S-PMSI and (C-\*, C-\*) S-PMSI.

### C-multicast Route Processing

C-multicast route processing functionality follows RFC 6513 section 8.1.2 (BGP used for route exchange). The processing is analogous to BGP Unicast VPN route exchange. [Figure 52](#) shows C-multicast route processing with non-segmented mLDP PMSI details.

**Figure 52 Non-segmented mLDP C-multicast Exchange (Option C)**



1030

## LEAF Node Cavities



**Caution:** The SR OS ASBR does not currently support receiving a non-recursive opaque FEC (opaque type 1).

The LEAF (PE-2) has to have the ROOT-1 system IP installed in RTM via BGP. If the ROOT-1 is installed in RTM via IGP, the LEAF will not generate the recursive opaque FEC. As such, the ASBR 3 will not process the LDP FEC correctly.

### 3.4.16.5.2 Configuration Example

No configuration is required for Option B or Option C on ASBRs, although for Option B, **config>router>bgp>enable-inter-as-vpn** is required to enable inter-as-non-segmented MLDP through the ASBR router.

Policy is required for a root or leaf PE for removing the NO\_EXPORT community from MVPN routes, which can be configured using an export policy on the PE.

The following is an example for configuring a policy on PEs to remove the **no-export**:

```
*A:Dut-A>config>router>policy-options# info

 community "no-export" members "no-export"
 policy-statement "remNoExport"
 default-action accept
 community remove "no-export"
 exit
 exit

*A:Dut-A>config>router>policy-options#
```

The following is an example for configuring in BGP the policy in a global, group, or peer context:

```
*A:Dut-A>config>router>bgp# info

 vpn-apply-export
 export "remNoExport"

```

### 3.4.16.5.3 Inter-AS Non-segmented MLDP

Refer to the “Inter-AS Non-segmented MLDP” section of the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR MPLS Guide* for more information.

#### 3.4.16.5.4 ECMP

Refer to the “ECMP” section of the *7450 ESS, 7750 SR, 7950 XRS, and VSR MPLS Guide* for more information about ECMP.

### 3.4.17 Weighted ECMP and ECMP for VPRN IPv4 and IPv6 over MPLS LSPs

ECMP over MPLS LSPs for VPRN services refers to spraying packets across multiple named RSVP and SR-TE LSPs within the same ECMP set.

The ECMP-like spraying consists of hashing the relevant fields in the header of a labeled packet and selecting the next-hop tunnel based on the modulo operation of the output of the hash and the number of ECMP tunnels. The maximum number of ECMP tunnels selected from the TTM matches the value of the user-configured **ecmp** option. Only LSPs with the same lowest LSP metric can be part of the ECMP set. If the number of these LSPs is higher than the value configured in the **ecmp** option, the LSPs with the lowest tunnel IDs are selected first.

In weighted ECMP, the load-balancing weight of the LSP is normalized by the system and then used to bias the amount of traffic forwarded over each LSP. The weight of the LSP is configured using the **config>router>mpls>lsp>load-balancing-weight weight** and **config>router>mpls>lsp-template>load-balancing-weight weight** commands.

If one or more LSPs in the ECMP set have no **load-balancing-weight** configured, and the ECMP is set to a specific next hop, regular ECMP spraying is used.

Weighted ECMP is configured for VPRN services with SDP auto-bind by using the **config>service>vprn>auto-bind-tunnel>ecmp max-ecmp-routes** and **config>service>vprn>auto-bind-tunnel>weighted-ecmp** commands. Weighted ECMP is disabled by default.

The **ecmp max-ecmp-routes** command allows explicit configuration of the number of tunnels that **auto-bind-tunnel** can use to resolve for a VPRN. The *max-ecmp-routes* parameter range is 0 to 32.

If weighted ECMP is enabled, then a path is selected based on the output of the hashing algorithm. Packet paths are then mapped to LSPs in the SDP in proportion to the configured load-balancing weight of the LSP. The hash is based on the system load-balancing configuration.

---

## 3.5 FIB Prioritization

The RIB processing of specific routes can be prioritized through the use of the **rib-priority** command. This command allows specific routes to be prioritized through the protocol processing so that updates are propagated to the FIB as quickly as possible.

The **rib-priority** command can be configured within the VPRN instance of the OSPF or IS-IS routing protocols. For OSPF, a prefix list can be specified that identifies which route prefixes should be considered high priority. If the **rib-priority high** command is configured under an **VPRN>OSPF>area>interface** context then all routes learned through that interface is considered high priority. For the IS-IS routing protocol, RIB prioritization can be either specified through a prefix-list or an IS-IS tag value. If a prefix list is specified then route prefixes matching any of the prefix list criteria will be considered high priority. If instead an IS-IS tag value is specified then any IS-IS route with that tag value will be considered high priority.

The routes that have been designated as high priority will be the first routes processed and then passed to the FIB update process so that the forwarding engine can be updated. All known high priority routes should be processed before the routing protocol moves on to other standard priority routes. This feature will have the most impact when there are a large number of routes being learned through the routing protocols.

## 3.6 Configuring a VPRN Service with CLI

This section provides information to configure Virtual Private Routed Network (VPRN) services using the command line interface.

### 3.6.1 Basic Configuration

The following fields require specific input (there are no defaults) to configure a basic VPRN service:

- customer ID (refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Services Overview Guide*)
- specify interface parameters
- specify spoke SDP parameters

The following example displays a sample configuration of a VPRN service.

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

vrf-import "vrfImpPolCust1"
vrf-export "vrfExpPolCust1"
ecmp 8
autonomous-system 10000
route-distinguisher 10001:1
auto-bind-tunnel
 resolution filter
 resolution-filter ldp
vrf-target target:10001:1
interface "to-cel" create
 address 10.1.0.1/24
 proxy-arp
 exit
 sap 1/1/10:1 create
 ingress
 qos 100
 exit
 egress
 qos 1010
 filter ip 10
 exit
 exit
 dhcp
 description "DHCP test"
 exit
 vrrp 1
 exit
exit
static-route-entry 10.5.0.0/24
 next-hop 10.1.1.2
```

```
bgp
 router-id 10.0.0.1
 group "to-cel"
 export "vprnBgpExpPolCust1"
 peer-as 65101
 neighbor 10.1.1.2
 exit
exit
exit
pim
 apply-to all
 rp
 static
 exit
 bsr-candidate
 shutdown
 exit
 rp-candidate
 shutdown
 exit
 exit
exit
rip
 export "vprnRipExpPolCust1"
 group "cel"
 neighbor "to-cel"
 exit
 exit
exit
no shutdown

*A:ALA-1>config>service>vprn#
```

## 3.6.2 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.

- Step 1.** Associate a VPRN service with a customer ID.
- Step 2.** Define an autonomous system (optional).
- Step 3.** Define a route distinguisher (mandatory).
- Step 4.** Define VRF route-target associations or VRF import/export policies.
- Step 5.** Define PIM parameters (optional).
- Step 6.** Create a subscriber interface (applies to the 7750 SR only and is optional).
- Step 7.** Create an interface.
- Step 8.** Define SAP parameters on the interface.
  - Select node(s) and port(s).



- Optional - select QoS policies other than the default (configured in config>qos context).
- Optional - select filter policies (configured in config>filter context).
- Optional - select accounting policy (configured in config>log context).
- Optional - configure DHCP features. (applies to the 7450 ESS and 7750 SR)

**Step 9.** Define BGP parameters (optional).

- BGP must be enabled in the **config>router>bgp** context.

**Step 10.** Define RIP parameters (optional).

**Step 11.** Define spoke SDP parameters (optional).

**Step 12.** Create confederation autonomous systems within an AS. (optional).

**Step 13.** Enable the service.

### 3.6.3 Configuring VPRN Components

This section provides VPRN configuration examples.

#### 3.6.3.1 Creating a VPRN Service

Use the following CLI syntax to create a VPRN service. A route distinguisher must be defined and the VPRN service must be administratively up in order for VPRN to be operationally active.

**CLI Syntax:**

```
config>service# vprn service-id [customer customer-id]
route-distinguisher [ip-address:number1 | asn:number2]
description description-string
no shutdown
```

The following example displays a VPRN service configuration.

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

...
 vprn 1 customer 1 create
 route-distinguisher 10001:0
 no shutdown
 exit
...

*A:ALA-1>config>service>vprn#
```

### 3.6.3.2 Configuring Global VPRN Parameters

Refer to [VPRN Service Configuration Commands](#) for CLI syntax to configure VPRN parameters.

The following example displays a VPRN service with configured parameters.

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

...
 vprn 1 customer 1 create
 vrf-import "vrfImpPolCust1"
 vrf-export "vrfExpPolCust1"
 autonomous-system 10000
 route-distinguisher 10001:1
 spoke-sdp 2 create
 exit
 no shutdown
 exit
...

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

### 3.6.3.3 Configuring VPRN Log Parameters

The following output displays a VPRN log configuration example.

```
B:Dut-C>config>service>vprn# info

 dhcp
 local-dhcp-server "vprn_1" create
 use-pool-from-client
 force-renews
 no shutdown
 exit
 exit
 snmp
 community "YsMv96H2KZVKQeakNAq.38gvyr.MH9vA" hash2 r version both
 community "gkYL94190FFgu91PiRNvn3Rnl0edkMU1" hash2 rw version v2c
 access
 log
 filter 1
 default-action forward
 entry 1
 action forward
 exit
 exit
 syslog 1
 address 3ffe::e01:403
 log-prefix "vprn1"
 exit
 snmp-trap-group 32001
```

```

 trap-target "3" address 3ffe::e01:403 port 9000 snmpv2c notify-
community "vprn1"
 exit
 log-id 1
 filter 1
 from main change
 to syslog 1
 exit
 log-id 3
 filter 1
 from main change
 to snmp
 exit
exit
...

B:Dut-C>config>service>vprn#

```

### 3.6.3.3.1 Configuring a Spoke-SDP

Use the following CLI syntax to configure spoke SDP parameters:

**CLI Syntax:**

```

config>service# vprn service-id [customer customer-id]
spoke-sdp sdp-id
 no shutdown
interface ip-int-name
 spoke-sdp sdp-id:vc-id [vc-type {ether | vlan}]
 egress
 filter {ip ip-filter-id}
 vc-label egress-vc-label
 ingress
 filter {ip ip-filter-id}
 vc-label ingress-vc-label
 tos-marking-state {trusted | untrusted}
 no shutdown

```

Use the following CLI syntax to configure spoke SDP parameters for the 7750 SR:

**CLI Syntax:**

```

config>service# vprn service-id [customer customer-id]
spoke-sdp sdp-id
 no shutdown
interface ip-int-name
 spoke-sdp sdp-id:vc-id [vc-type {ether | vlan |
vpls}]
 egress
 filter {ip ip-filter-id}
 vc-label egress-vc-label
 ingress
 filter {ip ip-filter-id}

```

```

 vc-label ingress-vc-label
 tos-marking-state {trusted | untrusted}
 no shutdown

```

The following output displays a spoke SDP configuration.

```

A:ALA-48>config>service>vprn# info

...
 interface "SpokeSDP" create
 spoke-sdp 3:4 create
 ingress
 vc-label 3000
 filter ip 10
 exit
 egress
 vc-label 2000
 filter ip 10
 exit
 exit
 exit
...
 spoke-sdp 3 create
 exit
 no shutdown

A:ALA-48>config>service>vprn#

```

### 3.6.3.4 Configuring VPRN Protocols - PIM

Refer to [VPRN Service Configuration Commands](#) for CLI syntax to configure VPRN parameters.

The following example displays a VPRN PIM configuration for the 7750 SR:

```

config>service# info
#-----
...
 vprn 1 customer 2 create
 route-distinguisher 1:11
 interface "if1" create
 address 10.13.14.15/32
 loopback
 exit
 interface "if2" create
 address 10.13.14.1/24
 sap 1/1/2:0 create
 exit
 exit
 pim
 interface "if1"
 exit

```

```

 interface "if2"
 exit
 rp
 static
 exit
 bsr-candidate
 shutdown
 exit
 rp-candidate
 shutdown
 exit
 exit
 exit
 no shutdown
 exit
exit
#-----
config>service#

```

### 3.6.3.4.1 Configuring Router Interfaces

Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide* for command descriptions and syntax information to configure router interfaces.

The following example displays a router interface configurations:

```

ALA48>config>router# info
#-----
echo "IP Configuration"
#-----
...
 interface "if1"
 address 10.2.2.1/24
 port 1/1/33
 exit
 interface "if2"
 address 10.49.1.46/24
 port 1/1/34
 exit
 interface "if3"
 address 10.11.11.1/24
 port 1/1/35
 exit
...
#-----
ALA48>config>router#

```

### 3.6.3.4.2 Configuring VPRN Protocols - BGP

The autonomous system number and router ID configured in the VPRN context only applies to that particular service.

The minimal parameters that should be configured for a VPRN BGP instance are:

- Specify an autonomous system number for the router. See [Configuring Global VPRN Parameters](#).
- Specify a router ID - 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 above peer.

VPRN BGP is administratively enabled upon creation. Minimally, to enable VPRN BGP in a VPRN instance, you must associate an autonomous system number and router ID for the VPRN service, create a peer group, neighbor, and associate a peer AS number. There are no default VPRN BGP groups or neighbors. Each VPRN BGP group and neighbor must be explicitly configured.

All parameters configured for VPRN BGP are applied to the group and are inherited by each peer, but a group parameter can be overridden on a specific basis. VPRN BGP command hierarchy consists of three levels:

- the global level
- the group level
- the neighbor level

For example:

**CLI Syntax:**

```
config>service>vprn>bgp#(global level)
 group (group level)
 neighbor (neighbor level)
```

The local-address must be explicitly configured if two systems have multiple BGP peer sessions between them for the session to be established.

For more information about the BGP protocol, refer to the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Router Configuration Guide*.

### Configuring VPRN BGP Group and Neighbor Parameters

A group is a collection of related VPRN BGP peers. The group name should be a descriptive name for the group. Follow your group, name, and ID naming conventions for consistency and to help when troubleshooting faults.

All parameters configured for a peer group are applied to the group and are inherited by each peer (neighbor), but a group parameter can be overridden on a specific neighbor-level basis.

After a group name is created and options are configured, neighbors can be added within the same autonomous system to create IBGP connections and/or neighbors in different autonomous systems to create EBGP peers. All parameters configured for the peer group level are applied to each neighbor, but a group parameter can be overridden on a specific neighbor basis.

### Configuring Route Reflection

Route reflection can be implemented in autonomous systems with a large internal BGP mesh to reduce the number of IBGP sessions required. One or more routers can be selected to act as focal points, for internal BGP sessions. Several BGP-speaking routers can peer with a route reflector. A route reflector forms peer connections to other route reflectors. A router assumes the role as a route reflector by configuring the **cluster** *cluster-id* command. No other command is required unless you want to disable reflection to specific peers.

If you configure the cluster command at the global level, then all subordinate groups and neighbors are members of the cluster. The route reflector cluster ID is expressed in dotted decimal notation. The ID should be a significant topology-specific value. No other command is required unless you want to disable reflection to specific peers.

If a route reflector client is fully meshed, the **disable-client-reflect** command can be enabled to stop the route reflector from reflecting redundant route updates to a client.

### Configuring BGP Confederations

A VPRN can be configured to belong to a BGP confederation. BGP confederations are one technique for reducing the degree of IBGP meshing within an AS. When the confederation command is in the configuration of a VPRN the type of BGP session formed with a VPRN BGP neighbor is determined as follows:

- The session is of type IBGP if the peer AS is the same as the local AS.
- The session is of type confed-EBGP if the peer AS is different than the local AS AND the peer AS is listed as one of the members in the confederation command.
- The session is of type EBGP if the peer AS is different than the local AS AND the peer AS is not listed as one of the members in the confederation command.

## VPRN BGP CLI Syntax

Use the CLI syntax to configure VPRN BGP parameters.

The following example displays a VPRN BGP configuration:

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

...
 vprn 1 customer 1 create
 vrf-import "vrfImpPolCust1"
 vrf-export "vrfExpPolCust1"
 ecmp 8
 autonomous-system 10000
 route-distinguisher 10001:1
 auto-bind-tunnel
 resolution filter
 resolution-filter ldp
 vrf-target target:10001:1
 interface "to-cel" create
 address 10.1.0.1/24
 sap 1/1/10:1 create
 ingress
 scheduler-policy "SLA2"
 qos 100
 exit
 egress
 scheduler-policy "SLA1"
 qos 1010
 filter ip 6
 exit
 exit
 exit
 static-route-entry 10.5.0.0/24
 next-hop 10.1.1.2
 bgp
 router-id 10.0.0.1
 group "to-cel"
 export "vprnBgpExpPolCust1"
 peer-as 65101
 neighbor 10.1.1.2
 exit
 exit
 spoke-sdp 2 create
 exit
 no shutdown
 exit
...

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



### 3.6.3.4.3 Configuring VPRN Protocols - RIP

PE routers which attach to a particular VPN need to know, for each of that VPN's sites, which addresses in that VPN are at each site. There are several ways that a PE router can obtain this set of addresses. The Routing Information Protocol (RIP) sends routing update messages that include entry changes. The routing table is updated to reflect the new information. This functionality applies only to the 7450 ESS and 7750 SR.

RIP can be used as a PE/CE distribution technique. PE and CE routers may be RIP peers, and the CE may use RIP to tell the PE router the set of address prefixes which are reachable at the CE router's site. When RIP is configured in the CE, care must be taken to ensure that address prefixes from other sites (i.e., address prefixes learned by the CE router from the PE router) are never advertised to the PE. Specifically, if a PE router receives a VPN-IPv4 route, and as a result distributes an IPv4 route to a CE, then that route must not be distributed back from that CE's site to a PE router (either the same router or different routers).

In order to enable a VPRN RIP instance, the RIP protocol must be enabled in the **config>service> >vprn>rip** context of the VPRN. VPRN RIP is administratively enabled upon creation. Configuring other RIP commands and parameters are optional.



**Caution:** Careful planning is essential to implement commands that can affect the behavior of VPRN RIP global, group, and neighbor levels. Because the RIP commands are hierarchical, analyze the values that can disable features on a particular level.

The parameters configured on the VPRN RIP global level are inherited by the group and neighbor levels. Many of the hierarchical VPRN RIP commands can be modified on different levels. The most specific value is used. That is, a VPRN RIP group-specific command takes precedence over a global VPRN RIP command. A neighbor-specific statement takes precedence over a global VPRN RIP and group-specific command. For example, if you modify a VPRN RIP neighbor-level command default, the new value takes precedence over VPRN RIP group- and global-level settings. There are no default VPRN RIP groups or neighbors. Each VPRN RIP group and neighbor must be explicitly configured.

The minimal parameters that should be configured for a VPRN instance are:

- Specify a VPRN RIP peer group.
- Specify a VPRN RIP neighbor with which to peer.
- Specify a VPRN RIP peer-AS that is associated with the above peer.

VPRN RIP command hierarchy consists of three levels:

- The global level
- The group level
- The neighbor level

For example:

**CLI Syntax:**   config>service>vprn>rip#(global level)  
                  group(group level)  
                  neighbor (neighbor level)

### VPRN RIP CLI Syntax

The following example displays a VPRN RIP configuration:

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

...
 vprn 1 customer 1 create
 vrf-import "vrfImpPolCust1"
 vrf-export "vrfExpPolCust1"
 ecmp 8
 autonomous-system 10000
 route-distinguisher 10001:1
 auto-bind-tunnel
 resolution filter
 resolution-filter ldp
 vrf-target target:10001:1
 interface "to-cel" create
 address 10.1.0.1/24
 sap 1/1/10:1 create
 ingress
 scheduler-policy "SLA2"
 qos 100
 exit
 egress
 scheduler-policy "SLA1"
 qos 1010
 filter ip 6
 exit
 exit
 exit
static-route-entry 10.5.0.0/24
 next-hop 10.1.1.2
bgp
 router-id 10.0.0.1
 group "to-cel"
 export "vprnBgpExpPolCust1"
 peer-as 65101
 neighbor 10.1.1.2
 exit
 exit
```

```

 exit
 rip
 export "vprnRipExpPolCust1"
 group "cel"
 neighbor "to-cel"
 exit
 exit
exit
spoke-sdp 2 create
exit
no shutdown
exit
...

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

```

For more information about the RIP protocol, refer to the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Router Configuration Guide*.

### 3.6.3.4.4 Configuring VPRN Protocols - OSPF

Each VPN routing instance is isolated from any other VPN routing instance, and from the routing used across the backbone. OSPF can be run with any VPRN, independently of the routing protocols used in other VPRNs, or in the backbone itself. For more information about the OSPF protocol, refer to the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Router Configuration Guide*.

**CLI Syntax:**    config>service>vprn>ospf#

#### VPRN OSPF CLI Syntax

Refer to [OSPF Commands](#) for CLI syntax to configure VPRN parameters.

The following example displays the VPRN OSPF configuration shown above:

```

*A:ALA-48>config>service# info

vprn 2 customer 1 create
 interface "test" create
 exit
 no shutdown
exit
 area 0.0.0.0
 virtual-link 1.2.3.4 transit-area 1.2.3.4
 hello-interval 9
 dead-interval 40
 exit
exit

*A:ALA-48>config>service#

```

For more information about the OSPF protocol, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide*.

### 3.6.3.4.5 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.

You can configure a VPRN interface as a loopback interface by issuing the loopback command instead of the **sap sap-id** command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

When using mtrace/mstat in a Layer 3 VPN context then the configuration for the VPRN should have a loopback address configured which has the same address as the core instance's system address (BGP next-hop).

Refer to [OSPF Commands](#) for CLI commands and syntax.

The following example displays a VPRN interface configuration:

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

...
 vprn 1 customer 1 create
 vrf-import "vrfImpPolCust1"
 vrf-export "vrfExpPolCust1"
 ecmp 8
 autonomous-system 10000
 route-distinguisher 10001:1
 auto-bind-tunnel
 resolution filter
 resolution-filter ldp
 vrf-target target:10001:1
 interface "to-cel" create
 address 10.1.0.1/24
 exit
 exit
 static-route-entry 10.5.0.0/24
 next-hop 10.1.1.2
 spoke-sdp 2 create
 exit
 no shutdown
 exit
 ...

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

### 3.6.3.4.6 Configuring Overload State on a Single SFM

When a router has fewer than the full set of SFMs functioning, the forwarding capacity can be reduced. Some scenarios include:

- fewer than the maximum number of SFMs installed in the system
- one or more SFMs have failed
- the system is in the ISSU process and the SFM is co-located on the CPM

An overload condition can be set for IS-IS and OSPF to enable the router to still participate in exchanging routing information, but route all traffic away from it when insufficient SFMs are active. This is achieved using the following CLI commands:

**CLI Syntax:**

```
config>router>single-sfm-overload [holdoff-time hold-off-time]
config>service>vprn>single-sfm-overload [holdoff-time hold-off-time]
tools>perform>redundancy>forced-single-sfm-overload
```

These cause an overload state in the IGP to trigger the traffic reroute by setting the overload bit in IS-IS or setting the metric to maximum in OSPF. When PIM uses IS-IS or OSPF to find out the upstream router, a next-hop change in the IS-IS or OSPF will cause PIM to join the new path and prune the old path, which effectively also reroutes the multicast traffic downstream as well as the unicast traffic.

When the problem is resolved, and the required compliment of SFMs become active in the router, the overload condition is cleared, which will cause the traffic to be routed back to the router.

The conditions to set overload are:

- 7950 XRS-20, 7750 SR-12/SR-7/SR-c12 and 7450 ESS-12/ESS-7/ESS-6 platforms: if an SF/CPM fails, the protocol will set the overload
- 7950-40 XRS and 7750 SR-12e platforms: if two SFMs fail (a connected pair on the XRS-40) the protocol will set the overload

### 3.6.3.4.7 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 SR. Each SAP must be unique within a router. A SAP cannot be defined if the interface **loopback** command is enabled.

When configuring VPRN interface SAP parameters, a default QoS policy is applied to each ingress and egress SAP. Additional QoS policies and scheduler policies must be configured in the **config>qos** context. Filter policies are configured in the **config>filter** context and must be explicitly applied to a SAP. There are no default filter policies.

VPRN interface ATM SAP parameters on a 7750 SR can only be configured on ATM-type MDAs and ATM-configured ports. The **periodic-loopback** command can only be enabled when the **config>system>atm>oam** context is enabled. See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide*.

Refer to [OSPF Commands](#) for CLI commands and syntax.

The following example displays a VPRN interface SAP configuration:

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

...
 vprn 1 customer 1 create
 vrf-import "vrfImpPolCust1"
 vrf-export "vrfExpPolCust1"
 ecmp 8
 autonomous-system 10000
 route-distinguisher 10001:1
 auto-bind-tunnel
 resolution filter
 resolution-filter ldp
 vrf-target target:10001:1
 interface "to-cel" create
 address 10.1.0.1/24
 sap 1/1/10:1 create
 ingress
 scheduler-policy "SLA2"
 qos 100
 exit
 egress
 scheduler-policy "SLA1"
 qos 1010
 filter ip 6
 exit
 exit
 exit
 static-route-entry 10.5.0.0/24
 next-hop 10.1.1.2
 spoke-sdp 2 create
 exit
 no shutdown
 exit
...

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

## 3.6.4 Configuring IPSec Parameters

The following output displays service with IPSec parameters configured.

```
*A:ALA-49>config# info

...
service
 ies 100 customer 1 create
 interface "ipsec-public" create
 address 10.10.10.1/24
 sap ipsec-1.public:1 create
 exit
 exit
 no shutdown
exit
vprn 200 customer 1 create
 ipsec
 security-policy 1 create
 entry 1 create
 local-ip 172.17.118.0/24
 remote-ip 172.16.91.0/24
 exit
 exit
 exit
 route-distinguisher 1:1
 interface "ipsec-private" tunnel create
 sap tunnel-1.private:1 create
 ipsec-tunnel "remote-office" create
 security-policy 1
 local-gateway-address 10.10.10.118 peer 10.10.7.91 delivery-
service 100
 dynamic-keying
 ike-policy 1
 pre-shared-key "humptydumpty"
 transform 1
 exit
 exit
 no shutdown
 exit
 exit
 exit
 interface "corporate-network" create
 address 172.17.118.118/24
 sap 1/1/2 create
 exit
 exit
 static-route-entry 172.16.91.0/24
 ipsec-tunnel "remote-office"
 no shutdown
 exit
 exit
...

*A:ALA-49>config#
```

## 3.7 Service Management Tasks

This section discusses VPRN service management tasks.

### 3.7.1 Modifying VPRN Service Parameters

Use the CLI syntax to modify VPRN parameters ([VPRN Service Configuration Commands](#)).

The following example displays the VPRN service creation output.

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

...
vprn 1 customer 1 create
 shutdown
 vrf-import "vrfImpPolCust1"
 vrf-export "vrfExpPolCust1"
 ecmp 8
 maximum-routes 2000
 autonomous-system 10000
 route-distinguisher 10001:1
 interface "to-cel" create
 address 10.1.1.1/24
 sap 1/1/10:1 create
 exit
 exit
 static-route-entry 10.5.0.0/24
 next-hop 10.1.1.2
 bgp
 router-id 10.0.0.1
 group "to-cel"
 export "vprnBgpExpPolCust1"
 peer-as 65101
 neighbor 10.1.1.2
 exit
 exit
 exit
 spoke-sdp 2 create
 exit
exit
...

*A:ALA-1>config>service>vprn#
```



## 3.7.2 Deleting a VPRN Service

A VPRN service cannot be deleted until SAPs and interfaces are shut down and deleted. If protocols and/or a spoke-SDP are defined, they must be shut down and removed from the configuration as well.

Use the following CLI syntax to delete a VPRN service:

**CLI Syntax:**

```
config>service#
[no] vprn service-id [customer customer-id]
shutdown
[no] interface ip-int-name
shutdown
[no] sap sap-id]
[no] bgp
shutdown
[no] rip
shutdown
[no] spoke-sdp sdp-id
[no] shutdown
```

## 3.7.3 Disabling a VPRN Service

A VPRN service can be shut down without deleting any service parameters.

**CLI Syntax:**

```
config>service#
vprn service-id [customer customer-id]
shutdown
```

**Example:**

```
config>service# vprn 1
config>service>vprn# shutdown
config>service>vprn# exit
```

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

...
vprn 1 customer 1 create
shutdown
vrf-import "vrfImpPolCust1"
vrf-export "vrfExpPolCust1"
ecmp 8
autonomous-system 10000
route-distinguisher 10001:1
auto-bind-tunnel
resolution filter
resolution-filter ldp
vrf-target target:10001:1
```

```

interface "to-cel" create
 address 10.1.0.1/24
 sap 1/1/10:1 create
 ingress
 scheduler-policy "SLA2"
 qos 100
 exit
 egress
 scheduler-policy "SLA1"
 qos 1010
 filter ip 6
 exit
 exit
exit
static-route-entry 10.5.0.0/24
 next-hop 10.1.1.2
bgp
 router-id 10.0.0.1
 group "to-cel"
 export "vprnBgpExpPolCust1"
 peer-as 65101
 neighbor 10.1.1.2
 exit
exit
rip
 export "vprnRipExpPolCust1"
 group "cel"
 neighbor "to-cel"
 exit
exit
spoke-sdp 2 create
exit
exit
...

*A:ALA-1>config>service#

```

### 3.7.4 Re-enabling a VPRN Service

To re-enable a VPRN service that was shut down.

**CLI Syntax:**

```

config>service#
vprn service-id [customer customer-id]
no shutdown

```

## 3.8 VPRN Service Configuration Commands

### 3.8.1 Command Hierarchies

- [VPRN Service Configuration Commands](#)
- [L2TP Commands](#)
- [DHCP Commands](#)
- [GSMP Commands](#)
- [IGMP Commands](#)
- [IPsec Configuration Commands](#)
- [Log Commands](#)
- [Remote Server Commands](#)
- [VPRN Management Commands](#)
- [Multicast VPN Commands](#)
- [VPRN Management Commands](#)
- [Redundant Interface Commands](#)
- [Router Advertisement Commands](#)
- [NTP Commands](#)
- [NAT Commands](#)
- [Subscriber Interface Commands](#)
- [Interface Commands](#)
- [Network Interface Commands](#)
- [Interface Spoke SDP Commands](#)
- [Interface VRRP Commands](#)
- [Interface SAP Commands](#)
- [Interface SAP Tunnel Commands](#)
- [Routed VPLS Commands](#)
- [Oper Group Commands](#)
- [Network Ingress Commands](#)
- [BGP Configuration Commands](#)
- [BGP Group Configuration Commands](#)
- [BGP Group Neighbor Configuration Commands](#)
- [IS-IS Configuration Commands](#)

- [OSPF Configuration Commands](#)
- [PIM Configuration Commands](#)
- [MSDP Configuration Commands](#)
- [MLD Configuration Commands](#)
- [RIP Configuration Commands](#)
- [RADIUS Commands](#)
- [Web Portal Protocol Configuration Commands](#)
- [AARP Interface Commands](#)

### 3.8.1.1 VPRN Service Configuration Commands

```

config
 — service
 — vprn service-id [name name] [customer customer-id] [create]
 — no vprn service-id
 — aggregate ip-prefix/ip-prefix-length [summary-only] [as-set] [aggregator as-
 number:ip-address] [black-hole [generate-icmp]] [community comm-id
 [comm-id]] [local-preference local-pref] [description description]
 — aggregate ip-prefix/ip-prefix-length [summary-only] [as-set] [aggregator as-
 number:ip-address] [community comm-id [comm-id]] [indirect ip-address]
 [local-preference local-pref] [description description]
 — no aggregate ip-prefix/ip-prefix-length
 — [no] allow-export-bgp-vpn
 — auto-bind-tunnel
 — ecmp max-ecmp-routes
 — no ecmp
 — [no] enforce-strict-tunnel-tagging
 — resolution {any | filter | disabled}
 — resolution-filter
 — [no] bgp
 — [no] gre
 — [no] ldp
 — [no] rib-api
 — [no] rsvp
 — [no] sr-isis
 — [no] sr-ospf
 — [no] sr-policy
 — [no] sr-te
 — [no] udp
 — [no] weighted-ecmp
 — autonomous-system as-number
 — no autonomous-system
 — backup-path [ipv4] [ipv6] [label-ipv4]
 — [no] carrier-carrier-vpn
 — confederation confed-as-num members as-number [as-number]
 — no confederation confed-as-num members as-number [as-number]
 — no confederation

```

---

```

— description description-string
— no description
— [no] dns
 — ipv4-source-address ipv4-address
 — no ipv4-source-address
 — ipv6-source-address ipv6-address
 — no ipv6-source-address
 — primary-dns ip-address
 — no primary-dns
 — secondary-dns ip-address
 — no secondary-dns
 — [no] shutdown
 — tertiary-dns ip-address
 — no tertiary-dns
— ecmp max-ecmp-routes
— no ecmp
— [no] entropy-label
— eth-cfm
 — tunnel-fault [accept | ignore]
— export-grt plcy-or-long-expr [plcy-or-expr [plcy-or-expr... (up to 4 max)]]
— no export-grt
— no export-inactive-bgp
— fib-priority {high | standard}
— flowspec
 — ip-filter-max-size {value | default}
 — ipv6-filter-max-size {value | default}
— grt-lookup
 — [no] enable-grt
 — [no] allow-local-management
 — export-grt plcy-or-long-expr [plcy-or-expr [plcy-or-expr]]
 — export-limit num-routes
 — no export-limit
 — export-v6-limit
 — no export-v6-limit
 — import-grt plcy-or-long-expr [plcy-or-expr [plcy-or-expr]]
 — no import-grt
— [no] hash-label
— igmp-host-tracking
 — expiry-time expiry-time
 — no expiry-time
 — [no] shutdown
— label-mode {vrf | next-hop}
— no label-mode
— maximum-ipv6-routes number [log-only] [threshold percent]
— no maximum-ipv6-routes
— maximum-routes number [log-only] [threshold percent]
— no maximum-routes
— mc-maximum-routes number [log-only] [threshold percent]
— no mc-maximum-routes
— multicast-info-policy policy-name
— no multicast-info-policy
— mvpn
— [no] nat
— network

```

```

— ingress
 — filter {ip ip-filter-id | ipv6 ipv6-filter-id}
 — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
 — qos network-policy-id fp-redirect-group queue-group-name
 instance instance-id
 — no qos
 — [no] urpf-check
— [no] ptp
 — peer a.b.c.d [create]
 — no peer a.b.c.d
 — local-priority local-priority
 — [no] log-sync-interval log-interval
 — [no] shutdown
 — peer-limit limit
 — no peer-limit
 — [no] shutdown
— reassembly-group nat-group-id
— route-distinguisher [ip-address:number1 | asn:number2 | auto-rd]
— no route-distinguisher
— router-id ip-address
— no router-id
— sgt-qos
 — application dscp-app-name dscp {dscp-value | dscp-name}
 — application dot1p-app-name dot1p dot1p-priority
 — no application {dscp-app-name | dot1p-app-name}
 — dscp dscp-name fc fc-name
 — no dscp dscp-name
— [no] shutdown
— single-sfm-overload [holdoff-time holdoff-time]
— no single-sfm-overload
— snmp
 — [no] access
 — community community-name [hash | hash2] [access-permissions]
 [version SNMP-version] [src-access-list list-name]
 — no community community-name [hash | hash2]
— source-address
 — application app [ip-int-name | ip-address]
 — no application app
 — application6 app ipv6-address
— [no] spoke-sdp sdp-id
 — [no] control-channel-status
 — [no] acknowledgment
 — refresh-timer value
 — no refresh-timer
 — request-timer timer1 retry-timer timer2 [timeout-multiplier
 multiplier]
 — no request-timer
 — [no] control-word
 — [no] pw-path-id
 — agi agi
 — no agi
 — saii-type2 global-id:node-id:ac-id
 — no saii-type2
 — taii-type2 global-id:node-id:ac-id

```

- 
- **no taii-type2**
  - **[no] shutdown**
  - **[no] static-route-entry** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**mcast**]
    - **[no] black-hole**
      - **[no] community** *comm-id*
      - **[no] description** *description-string*
      - **[no] generate-icmp**
      - **[no] metric** *metric-value*
      - **[no] preference** *preference-value*
      - **[no] prefix-list** *name* {**all** | **none** | **any**}
      - **[no] shutdown**
      - **[no] tag** *tag-value*
    - **community** *comm-id* [*comm-id*]
    - **no community**
    - **[no] grt**
      - **[no] description** *description-string*
      - **[no] metric** *metric-value*
      - **[no] preference** *preference-value*
      - **[no] shutdown**
    - **[no] indirect** *ip-address*
      - **[no] community** *comm-id*
      - **[no] cpe-check** *cpe-ip-address*
        - **[no] drop-count** *count*
        - **[no] interval** *seconds*
        - **[no] log**
        - **[no] padding-size** *padding-size*
      - **[no] description** *description-string*
      - **[no] destination-class** *dest-index*
      - **[no] forwarding-class**
        - **[no] priority** {**low** | **high**}
      - **[no] metric** *metric-value*
      - **[no] preference** *preference-value*
      - **[no] prefix-list** *name* {**all** | **none** | **any**}
      - **[no] shutdown**
      - **[no] source-class** *source-index*
      - **[no] tag** *tag-value*
    - **[no] ipsec-tunnel**
      - **[no] community** *comm-id*
      - **[no] description** *description-string*
      - **[no] destination-class** *dest-index*
      - **[no] forwarding-class** {**be** | **l2** | **af** | **l1** | **h2** | **ef** | **h1** | **nc**}
        - **[no] priority** {**low** | **high**}
      - **[no] metric** *metric-value*
      - **[no] preference** *preference-value*
      - **[no] shutdown**
      - **[no] source-class** *source-index*
      - **[no] tag** *tag-value*
    - **[no] next-hop** {*ip-address* | *ip-int-name* | *ipv6 address*}
      - **[no] bfd-enable**
      - **[no] community** *comm-id*
      - **[no] cpe-check** *cpe-ip-address*
        - **[no] drop-count** *count*
        - **[no] interval** *seconds*
        - **[no] log**

- [no] **padding-size** *padding-size*
- [no] **description** *description-string*
- [no] **destination-class** *dest-index*
- [no] **forwarding-class**
  - [no] **priority** {*low* | *high*}
- **load-balancing-weight** *value*
- **no load-balancing-weight**
- [no] **metric** *metric-value*
- [no] **preference** *preference-value*
- [no] **prefix-list** *name* {*all* | *none* | *any*}
- [no] **shutdown**
- [no] **source-class** *source-index*
- [no] **tag** *tag-value*
- [no] **validate-next-hop**
- **tag** *tag*
- **no tag**
- **ttl-propagate**
  - **local** [*inherit* | *none* | *vc-only* | *all*]
  - **transit** [*inherit* | *none* | *vc-only* | *all*]
- **type** {*hub* | *spoke* | *subscriber-split-horizon*}
- **no type**
- **vrf-export** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]
- **no vrf-export**
- **vrf-import** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]
- **no vrf-import**
- **vrf-target** {*ext-comm* | {[**export** *ext-comm*] [**import** *ext-comm*]}}
- **no vrf-target**
- [no] **weighted-ecmp**

### 3.8.1.2 L2TP Commands

- ```

config
  — service
    — vprn service-id [customer customer-id]
    — no vprn service-id
      — l2tp
      — no l2tp
        — avp-hiding {sensitive | always}
        — no avp-hiding
        — calling-number-format ascii-spec
        — no calling-number-format
        — challenge {always}
        — no challenge
        — destruct-timeout destruct-timeout
        — no destruct-timeout
        — exclude-avps calling-number
        — no exclude-avps
        — group tunnel-group-name [create]
        — no group tunnel-group-name
          — avp-hiding {sensitive | always | never}
          — no avp-hiding

```


-
- **challenge** {always | never}
 - **no challenge**
 - **description** description-string
 - **no description**
 - **destruct-timeout** destruct-timeout
 - **no destruct-timeout**
 - **hello-interval** hello-interval
 - **no hello-interval**
 - **idle-timeout** idle-timeout
 - **no ipcp-subnet-negotiation**
 - **no idle-timeout**
 - **l2tpv3**
 - **cookie-length** {4 | 8 | default}
 - **no cookie-length**
 - **digest-type** {default | md5 | sha1 | none}
 - **no digest-type**
 - **nonce-length** {length | default}
 - **no nonce-length**
 - **password** password [hash | hash2]
 - **no password**
 - **pw-cap-list** {ethernet | ethernet-vlan}
 - **no pw-cap-list**
 - **rem-router-id** ip-addr
 - **no rem-router-id**
 - **private-tcp-mss-adjust** octets
 - **no private-tcp-mss-adjust**
 - **public-tcp-mss-adjust** octets
 - **no public-tcp-mss-adjust**
 - [no] **track-password-change**
 - **lns-group** lns-group-id
 - **no lns-group**
 - **local-address** ip-address
 - **no local-address**
 - **local-name** host-name
 - **no local-name**
 - **max-retries-estab** max-retries
 - **no max-retries-estab**
 - **max-retries-not-estab** max-retries
 - **no max-retries-not-estab**
 - **password** password [hash | hash2]
 - **no password**
 - **ppp**
 - **authentication** {chap | pap | pref-chap | pref-pap}
 - **authentication-policy** auth-policy-name
 - **no authentication-policy**
 - **default-group-interface** ip-int-name service-id service-id
 - **default-group-interface** ip-int-name service-name svc-name
 - **no default-group-interface**
 - [no] **ipcp-subnet-negotiation**
 - **keepalive** seconds [hold-up-multiplier multiplier]
 - **no keepalive**
 - [no] **lcp-force-ack-accm**

-
- **mtu** *mtu-bytes*
 - **no mtu**
 - **[no] proxy-authentication**
 - **[no] proxy-lcp**
 - **user-db** *local-user-db-name*
 - **no user-db**
 - **session-assign-method** {existing-first | weighted | weighted-random}
 - **no session-assign-method**
 - **session-limit** *session-limit*
 - **session-limit unlimited**
 - **no session-limit**
 - **tunnel** *tunnel-name* [create]
 - **no tunnel** *tunnel-name*
 - **[no] auto-establish**
 - **avp-hiding** {never | sensitive | always}
 - **no avp-hiding**
 - **challenge** {always | never}
 - **no challenge**
 - **description** *description-string*
 - **no description**
 - **destruct-timeout** *destruct-timeout*
 - **no destruct-timeout**
 - **hello-interval** *hello-interval*
 - **hello-interval infinite**
 - **no hello-interval**
 - **idle-timeout** *idle-timeout*
 - **idle-timeout infinite**
 - **no idle-timeout**
 - **local-address** *ip-address*
 - **no local-address**
 - **local-name** *host-name*
 - **no local-name**
 - **max-retries-estab** *max-retries*
 - **no max-retries-estab**
 - **max-retries-not-estab** *max-retries*
 - **no max-retries-not-estab**
 - **password** *password* [hash | hash2]
 - **no password**
 - **peer** *ip-address*
 - **no peer**
 - **ppp**
 - **[no] ipcp-subnet-negotiation**
 - **[no] lcp-force-ack-accm**
 - **preference** *preference*
 - **no preference**
 - **remote-name** *host-name*
 - **no remote-name**
 - **session-limit** *session-limit*
 - **session-limit unlimited**
 - **no session-limit**
 - **[no] shutdown**
 - **l2tpv3**
 - **cookie-length** {4 | 8}

- **no cookie-length**
- **digest-type** {md5 | sha1 | none}
- **no digest-type**
- **nonce-length** *length*
- **no nonce-length**
- **password** *password* [hash | hash2]
- **no password**
- **private-tcp-mss-adjust** *octets*
- **no private-tcp-mss-adjust**
- **public-tcp-mss-adjust** **auto**
- **public-tcp-mss-adjust** *octets*
- **no public-tcp-mss-adjust**
- **transport-type** **ip**
- **no transport-type**
- **peer-address-change-policy** {accept | ignore | reject}
- **receive-window-size** *window-size*
- **no receive-window-size**
- **rtm-debounce-time** *debounce-time*
- **no rtm-debounce-time**
- **session-limit** *session-limit*
- **session-limit** **unlimited**
- **no session-limit**
- **[no] shutdown**

3.8.1.3 DHCP 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
 — failover
 — maximum-client-lead-time [hrs hours] [min minutes] [sec seconds]
 — no maximum-client-lead-time
 — partner-down-delay [hrs hours] [min minutes] [sec seconds]
 — no partner-down-delay
 — peer ip-address tag sync-tag-name
 — no peer ip-address
 — [no] shutdown
 — [no] startup-wait-time [min minutes] [sec seconds]
 — [no] force-renews
 — pool pool-name [create]
 — no pool pool-name
 — description description-string
 — no description

```

- 
- **failover**
    - **maximum-client-lead-time** [hrs hours] [min minutes] [sec seconds]
    - **no maximum-client-lead-time**
    - **partner-down-delay** [hrs hours] [min minutes] [sec seconds]
    - **no partner-down-delay**
    - **peer** ip-address tag sync-tag-name
    - **no peer** ip-address
    - **[no] shutdown**
    - **[no] startup-wait-time** [min minutes] [sec seconds]
  - **max-lease-time** [days days] [hrs hours] [min minutes] [sec seconds]
  - **no max-lease-time**
  - **min-lease-time** [days days] [hrs hours] [min minutes] [sec seconds]
  - **no min-lease-time**
  - **minimum-free** minimum-free [percent] [event-when-depleted]
  - **no minimum-free**
  - **offer-time** [min minutes] [sec seconds]
  - **no offer-time**
  - **options**
    - **custom-option** option-number address [ip-address]
    - **custom-option** option-number hex hex-string
    - **custom-option** option-number string ascii-string
    - **no custom-option** option-number
    - **dns-server** ip-address [ip-address]
    - **domain-name** domain-name
    - **no domain-name**
    - **lease-rebind-time** [days days] [hrs hours] [min minutes] [sec seconds]
    - **no lease-rebind-time**
    - **lease-renew-time** [days days] [hrs hours] [min minutes] [sec seconds]
    - **no lease-renew-time**
    - **lease-time** [days days] [hrs hours] [min minutes] [sec seconds]
    - **no lease-time**
    - **netbios-name-server** ip-address [ip-address]
    - **no netbios-name-server**
    - **netbios-node-type** netbios-node-type
    - **no netbios-node-type**
  - **subnet** {ip-address/mask | ip-address netmask} [create]
  - **no subnet** {ip-address/mask | ip-address netmask}
    - **[no] address-range** start-ip-address end-ip-address
    - **[no] drain**
    - **[no] exclude-addresses** start-ip-address [end-ip-address]
    - **maximum-declined** maximum-declined
    - **no maximum-declined**

- **minimum-free** *minimum-free* [**percent**] [**event-when-depleted**]
- **no minimum-free**
- **options**
  - **custom-option** *option-number* **address** [*ip-address*]
  - **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*]
  - **no default-router**
  - **subnet-mask** *ip-address*
  - **no subnet-mask**
- **[no] shutdown**
- **[no] use-gi-address**
- **[no] use-pool-from-client**
- **user-db** *local-user-db-name*
- **no user-db**
- **dhcp6**
  - **local-dhcp-server** *server-name* [**create**]
  - **no local-dhcp-server** *server-name*
    - **description** *description-string*
    - **no description**
    - **failover**
      - **ignore-mclt-on-takeover**
      - **no ignore-mclt-on-takeover**
      - **maximum-client-lead-time** [*hrs hours*] [*min minutes*] [*sec seconds*]
      - **no maximum-client-lead-time**
      - **partner-down-delay** [*hrs hours*] [*min minutes*] [*sec seconds*]
      - **no partner-down-delay**
      - **peer** *ip-address* **tag** *sync-tag-name*
      - **no peer** *ip-address*
      - **[no] shutdown**
      - **[no] startup-wait-time** [*min minutes*] [*sec seconds*]
  - **[no] ignore-rapid-commit**
  - **lease-hold-time** [*days days*] [*hrs hours*] [*min minutes*] [*sec seconds*]
  - **no lease-hold-time**
  - **pool** *pool-name* [**create**]
  - **no pool** *pool-name*
    - **description** *description-string*
    - **no description**
    - **options**
      - **custom-option** *option-number* **address** [*ipv6-address*]
      - **custom-option** *option-number* **hex** *hex-string*
      - **custom-option** *option-number* **string** *ascii-string*
      - **no custom-option** *option-number*
      - **delegated-prefix-length** *prefix-length*
      - **dns-server** *ipv6-address* [*ipv6-address*]
      - **domain-name** *domain-name*
      - **no domain-name**

- **prefix** *ipv6-address/prefix-length* [**failover** {**local** | **remote**}] [**pd**] [**wan-host**] [**create**]
- **no prefix** *ipv6-address/prefix-length*
  - **preferred-lifetime** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
  - **no preferred-lifetime**
  - **rebind-timer** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
  - **no rebind-timer**
  - **renew-timer** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
  - **no renew-timer**
  - **valid-lifetime** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
  - **no valid-lifetime**
- **use-link-address** [**scope** *scope*]
- **no use-link-address**
- [**no**] **use-pool-from-client**
- **user-ident** *user-ident*
- **no user-ident**

### 3.8.1.4 GSMP Commands

- ```

config
  — service
    — vpn service-id [customer customer-id]
    — no vpn service-id
      — gsmp
        — [no] group name
          — anccp
            — [no] dynamic-topology-discover
            — [no] oam
          — description description-string
          — no description
          — hold-multiplier multiplier
          — no hold-multiplier
          — keepalive seconds
          — no keepalive
          — [no] neighbor ip-address
            — description description-string
            — no description
            — local-address ip-address
            — no local-address
            — priority-marking dscp dscp-name
            — priority-marking prec ip-prec-value
            — no priority-marking
            — [no] shutdown
          — idle-filter
          — [no] idle-filter
          — persistence-database
          — [no] persistence-database

```

- [no] **shutdown**
- [no] **shutdown**

3.8.1.5 IGMP Commands

```

config
  — service
    — vrpn service-id [customer customer-id]
    — no vrpn service-id
    — [no] igmp
      — [no] group-interface ip-int-name
      — [no] group-interface fwd-service service-id ip-int-name
        — [no] disable-router-alert-check
        — import policy-name
        — no import
        — max-groups value
        — no max-groups
        — max-grp-sources max-group-sources
        — no max-grp-sources
        — max-sources max-sources
        — no max-sources
      — mcac
        — if-policy mcac-if-policy-name
        — no if-policy
        — mc-constraints
          — [no] shutdown
        — policy policy-name
        — no policy
        — unconstrained-bw bandwidth mandatory-bw
          mandatory-bw
        — no unconstrained-bw
      — query-src-ip ip-address
      — no query-src-ip
      — [no] shutdown
      — [no] sub-hosts-only
      — [no] subnet-check
      — version version
      — no version
    — grp-if-query-src-ip ip-address
    — no grp-if-query-src-ip
    — [no] interface ip-int-name
      — [no] disable-router-alert-check
      — import policy-name
      — no import
      — max-groups value
      — no max-groups
      — max-sources max-sources
      — no max-sources
      — max-grp-sources max-grp-sources
      — no max-grp-sources
      — mcac

```

- **if-policy** *mcac-if-policy-name*
- **no if-policy**
- **mc-constraints**
 - **level** *level-id* **bw** *bandwidth*
 - **no level** *level-id*
 - **number-down** *number-lag-port-down* **level** *level-id*
 - **no number-down** *number-lag-port-down*
 - **[no] shutdown**
 - **[no] use-lag-port-weight**
- **policy** *policy-name*
- **no policy**
- **unconstrained-bw** *bandwidth* **mandatory-bw** *mandatory-bw*
- **no unconstrained-bw**
- **[no] shutdown**
- **ssm-translate**
 - **[no] grp-range** *start end*
 - **[no] source** *ip-address*
- **static**
 - **[no] group**
 - **[no] source** *ip-address*
 - **[no] starg**
 - **[no] group** *grp-ip-address*
 - **[no] group start** *grp-ipv6-address* **end** *grp-ipv6-address* **[step** *ipv6-address* **]**
- **[no subnet-check]**
- **version** *version*
- **no version**
- **[no] query-interval**
- **query-interval** *seconds*
- **[no] query-last-member-interval**
- **query-last-member-interval** *seconds*
- **[no] query-response-interval**
- **query-response-interval** *seconds*
- **[no] robust-count**
- **robust-count** *robust-count*
- **[no] shutdown**
- **ssm-translate**
 - **[no] grp-range** *start end*
 - **[no] source** *ip-address*

3.8.1.6 IPsec Configuration Commands

- ```

config
 — service
 — vprn service-id [customer customer-id]
 — no vprn service-id
 — ipsec
 — security-policy security-policy-id [create]
 — no security-policy security-policy-id

```



---

```

— entry entry-id [create]
— no entry entry-id
 — local-ip {ip-prefix/prefix-length | ip-prefix netmask | any}
 — remote-ip {ip-prefix/prefix-length | ip-prefix netmask |
 any}
— interface ip-int-name [create]
— no interface ip-int-name
 — [no] address {ip-address/mask | ip-address netmask}
 — description description-string
 — no description
 — ip-mtu octets
 — no ip-mtu
 — sap sap-id [create]
 — no sap sap-id
 — description description-string
 — no description
 — egress
 — [no] agg-rate
 — [no] limit-unused-bandwidth
 — [no] queue-frame-based-accounting
 — rate {max | rate}
 — no rate
 — filter ip ip-filter-id
 — no filter [ip ip-filter-id]
 — [no] qinq-mark-top-only
 — qos policy-id [port-redirect-group queue-group-name
 instance instance-id]
 — no qos
 — [no] queue-override
 — [no] queue queue-id
 — adaptation-rule [pir adaptation-rule] [cir
 adaptation-rule]
 — no adaptation-rule
 — avg-frame-overhead percentage
 — no avg-frame-overhead
 — cbs size-in-kbytes
 — no cbs
 — drop-tail low
 — percent-reduction-from-mbs
 percent
 — no percent-reduction-from-mbs
 — mbs {size [bytes | kilobytes] | default}
 — no mbs
 — parent [weight weight] [cir-weight cir-
 weight]
 — no parent
 — percent-rate pir-percent [cir cir-percent]
 — no percent-rate
 — rate pir-rate [cir cir-rate]
 — no rate
 — [no] scheduler-override
 — [no] scheduler scheduler-name
 — rate pir-rate [cir cir-rate]
 — no rate

```

- **scheduler-policy** *scheduler-policy-name*
- **no scheduler-policy**
- **ingress**
  - **filter ip** *ip-filter-id*
  - **no filter** [*ip ip-filter-id*]
  - **match-qinq-dot1p** {*top* | *bottom*}
  - **qos** *policy-id* [*shared-queuing* | *multipoint-shared*] [*fp-redirect-group queue-group-name instance instance-id*]
  - **no qos**
  - [no] **queue-override**
    - [no] **queue** *queue-id*
      - **adaptation-rule** [*pir adaptation-rule*] [*cir adaptation-rule*]
      - **no adaptation-rule**
      - **cbs** *size-in-kbytes*
      - **no cbs**
      - **drop-tail low**
        - **percent-reduction-from-mbs** *percent*
        - **no percent-reduction-from-mbs**
      - **mbs** {*size* [*bytes* | *kilobytes*] | *default*}
      - **no mbs**
      - **rate** *pir-rate* [*cir cir-rate*]
      - **no rate**
  - [no] **scheduler-override**
    - [no] **scheduler** *scheduler-name*
      - **rate** *pir-rate* [*cir cir-rate*]
      - **no rate**
  - **scheduler-policy** *scheduler-policy-name*
  - **no scheduler-policy**
  - [no] **shutdown**
  - **ipsec-tunnel** *ipsec-tunnel-name* [*create*]
  - **no ipsec-tunnel** *ipsec-tunnel-name*
    - [no] **bfd-designate**
    - **bfd-enable service** *service-id interface interface-name dst-ip ip-address*
    - **bfd-enable service-name** *service-name interface interface-name dst-ip ip-address*
    - **no bfd-enable**
    - [no] **clear-df-bit**
    - **description** *description-string*
    - **no description**
    - [no] **dynamic-keying**
      - [no] **auto-establish**
      - **ike-policy** *ike-policy-id*
      - **no ike-policy**
      - **local-id type** {*ipv4 v4address* | *fqdn fqdn-value*}
      - **no local-id**
      - **pre-shared-key** *key*
      - **no pre-shared-key**
      - **transform** *transform-id* [*transform-id*]
      - **no transform**

- **local-gateway-address** *ip-address* **peer** *ip-address*  
**delivery-service** *service-id*
- **local-gateway-address** *ip-address* **peer** *ip-address*  
**delivery-service-name** *service-name*
- **no local-gateway-address**
- **[no] manual-keying**
  - **security-association** *security-entry-id*  
**authentication-key** *authentication-key* **encryption-key** *encryption-key*  
**spi** *spi* **transform** *transform-id*  
**direction** {inbound | outbound}
  - **no security-association** *security-entry-id*  
**direction** {inbound | outbound}
- **replay-window** {32 | 64 | 128 | 256 | 512}
- **no replay-window**
- **security-policy** *security-policy-id* [strict-match]
- **no security-policy**
- **[no] shutdown**
- **[no] sap** *sap-id*
  - **ipsec-gw** *name*
  - **no ipsec-gw**
    - **cert**
      - **cert-profile** *name*
      - **no cert-profile**
      - **trust-anchor-profile** *name*
      - **no trust-anchor-profile**
    - **default-secure-service** *service-id* **interface** *ip-int-name*
    - **default-secure-service** *name* *service-name* **interface**  
*ip-int-name*
    - **no default-secure-service**
    - **default-tunnel-template** *ipsec template identifier*
    - **no default-tunnel-template**
    - **ike-policy** *ike-policy-id*
    - **no ike-policy**
    - **local-gateway-address** *ip-address*
    - **no local-gateway-address**
    - **local-id** **type** {ipv4 | fqdn} [**value** *value*]
    - **no local-id**
    - **pre-shared-key** *key*
    - **no pre-shared-key**
    - **radius-accounting-policy** *policy-name*
    - **no radius-accounting-policy**
    - **radius-authentication-policy** *name*
    - **no radius-authentication-policy**
    - **[no] shutdown**

### 3.8.1.7 Log Commands

Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR System Management Guide* for information about configuring event and accounting logs.

```

config
 — service
 — vprn service-id [customer customer-id]
 — no vprn service-id
 — log
 — [no] filter filter-id
 — default-action {drop | forward}
 — no default-action
 — description description-string
 — no description
 — [no] entry entry-id
 — action {drop | forward}
 — no action
 — description description-string
 — no description
 — [no] match
 — application {eq | neq} application-id
 — no application
 — message {eq | neq} pattern pattern [regex]
 — no message
 — number {eq | neq | lt | lte | gt | gte} event-id
 — no number
 — severity {eq | neq | lt | lte | gt | gte} severity-level
 — no severity
 — subject {eq | neq} subject [regex]
 — no subject
 — [no] log-id log-id
 — description description-string
 — no description
 — filter filter-id
 — no filter
 — from {[main] | [change]}
 — no from
 — [no] shutdown
 — time-format {local | utc}
 — to snmp [size]
 — to syslog syslog-id
 — [no] snmp-trap-group log-id
 — description description-string
 — no description
 — trap-target name [address ip-address] [port port] [snmpv1 |
 snmpv2c | snmpv3] notify-community
 communityName | snmpv3SecurityName [security-level
 {no-auth-no-privacy | auth-no-privacy | privacy}]
 [replay]
 — no trap-target name
 — [no] syslog syslog-id
 — address ip-address
 — no address
 — description description-string
 — no description
 — facility syslog-facility
 — no facility

```

- **level** {emergency | alert | critical | error | warning | notice | info | debug}
- **no level**
- **log-prefix** *log-prefix-string*
- **no log-prefix**
- **port** *port*
- **no port**

### 3.8.1.8 Remote Server Commands

```

config
— service
 — vpn service-id [customer customer-id]
 — no vpn service-id
 — aaa
 — remote-servers
 — radius [create]
 — no radius
 — access-algorithm {direct | round-robin}
 — no access-algorithm
 — [no] accounting
 — accounting-port port
 — no accounting-port
 — [no] authorization
 — [no] interactive-authentication
 — port port
 — no port
 — retry count
 — no retry
 — server server-index address ip-address secret key [{hash | hash2}]
 — no server server-index
 — [no] shutdown
 — timeout seconds
 — no timeout
 — [no] use-default-template
 — tacplus [create]
 — no tacplus
 — accounting [record-type {start-stop | stop-only}]
 — no accounting
 — authorization [use-priv-lvl]
 — no authorization
 — [no] interactive-authentication
 — [no] priv-lvl-map
 — priv-lvl priv-lvl user-profile-name
 — no priv-lvl priv-lvl
 — server index address ip-address secret key [{hash | hash2}] [port port]
 — no server index
 — [no] shutdown
 — timeout seconds

```

- no **timeout**
- [no] **use-default-template**

### 3.8.1.9 VPRN Management Commands

```

config
 — service
 — vprn service-id [customer customer-id] [create]
 — no vprn service-id
 — management [create]
 — no management
 — [no] allow-ftp
 — [no] allow-ssh
 — [no] allow-telnet
 — [no] allow-telnet6

```

### 3.8.1.10 Multicast VPN Commands

```

config
 — service
 — vprn service-id [customer customer-id]
 — no vprn service-id
 — mvpn
 — [no] auto-discovery [default | mdt-safi] [source-address ip-address]
 — c-mcast-signaling {bgp | pim}
 — no c-mcast-signaling
 — intersite-shared [persistent-type5-adv] [kat-type5-adv-withdraw]
 — no intersite-shared
 — mdt-type {sender-receiver | sender-only | receiver-only}
 — red-source-list
 — src-prefix ip-address/mask [ip-address/mask]
 — no src-prefix ip-address/mask
 — ipv6
 — src-prefix ipv6-ip-address/prefix-length [ipv6-ip-address/prefix-length]
 — no src-prefix ipv6-ip-address/prefix-length
 — rpf-select
 — [no] core-mvpn service-id
 — group-prefix ip-address/mask [ip-address/mask] [starg]
 — no group-prefix ip-address/mask
 — provider-tunnel
 — inclusive
 — [no] bier
 — [no] shutdown
 — sub-domain sub-domain
 — no sub-domain
 — bsr {unicast | spmsi}
 — no bsr

```

- **mldp**
  - [no] **shutdown**
- **pim** {asm | ssm} *grp-ip-address*
- **no pim**
  - **hello-interval** *hello-interval*
  - **no hello-interval**
  - **hello-multiplier** *deci-units*
  - **no hello-multiplier**
  - [no] **improved-assert**
  - [no] **shutdown**
  - [no] **three-way-hello**
  - [no] **tracking-support**
- **rsvp**
  - [no] **enable-bfd-leaf**
  - **enable-bfd-root** [*transmit-interval*] [**multiplier** *multiplier*]
  - **no enable-bfd-root**
  - **lsp-template** *lsp-template*
  - **no lsp-template**
  - [no] **shutdown**
- [no] **wildcard-spmsi**
- **selective**
  - [no] **auto-discovery-disable**
  - [no] **bier**
    - [no] **shutdown**
    - **sub-domain** *sub-domain*
    - **no sub-domain**
- **data-delay-interval** *value*
- **no data-delay-interval**
- **data-threshold** {*c-grp-ip-addr/mask* | *c-grp-ip-addr netmask*} *s-pmsi-threshold* [**pe-threshold-add** *pe-threshold-add*] [**pe-threshold-delete** *pe-threshold-delete*]
- **data-threshold** *c-grp-ipv6-addr/prefix-length* *s-pmsi-threshold* [**pe-threshold-add** *pe-threshold-add*] [**pe-threshold-delete** *pe-threshold-delete*]
- **no data-threshold** {*c-grp-ip-addr/mask* | *c-grp-ip-addr netmask*}
- **no data-threshold** *c-grp-ipv6-addr/prefix-length*
- [no] **enable-asm-mdt**
- [no] **join-tlv-packing-disable**
- [no] **multistream-spmsi** *index* [**create**]
  - [no] **group** *ip-address* [*/mask*]
    - [no] **source** *ip-address* [*/mask*]
    - [no] **source any**
  - [no] **lsp-template** *lsp-template-name*
  - **mdt-pim** **mode** {asm | ssm} **group-address** *group-ip-address*
  - **no mdt-pim**
  - [no] **shutdown**
- [no] **maximum-p2mp-spmsi**
- [no] **pim-asm** {*grp-ip-address/mask* | *grp-ip-address netmask*}

- **pim-ssm** {*grp-ip-address/mask* | *grp-ip-address netmask*}
- **no rsvp**
  - **no lsp-template**
  - **no shutdown**
- **no mldp**
  - **no shutdown**
- **no pim-asm**
- **umh-pe-backup**
  - **umh-pe** *ip-address standby ip-address*
  - **no umh-pe** *ip-address*
- **umh-selection** {*highest-ip* | *hash-based* | *tunnel-status* | *unicast-rt-pref*}
- **no umh-selection**
- **vrf-export** {*unicast* | *policy-name [policy-name]*}
- **no vrf-export**
- **vrf-import** {*unicast* | *plcy-or-long-expr [plcy-or-expr [plcy-or-expr]]*}
- **no vrf-import**
- **vrf-target** {*unicast* | *ext-community* | **export** *unicast* | *ext-community* | **import** *unicast* | *ext-community*}
- **no vrf-target**
  - **export** {*unicast* | *ext-community*}
  - **import** {*unicast* | *ext-community*}

### 3.8.1.11 Redundant Interface Commands

- ```

config
  — service
    — vprn service-id [customer customer-id]
    — no vprn service-id
    — [no] redundant-interface ip-int-name
      — address {ip-address/mask | ip-address netmask} [remote-ip ip-address]
      — no address
      — description long-description-string
      — no description
      — hold-time
        — up ip seconds
        — no up ip
        — up ipv6 seconds
        — no up ipv6
        — down ip seconds [init-only]
        — no down ip
        — down ipv6 seconds [init-only]
        — no down ipv6
      — [no] shutdown
      — [no] spoke-sdp sdp-id:vc-id
        — egress
          — filter [ip ip-filter-id]
          — vc-label ingress-vc-label
          — no vc-label [ingress-vc-label]

```


- **ingress**
 - **filter** [*ip ip-filter-id*]
 - **no filter**
 - **vc-label** *ingress-vc-label*
 - **no vc-label** [*ingress-vc-label*]
- [no] **shutdown**

3.8.1.12 Router Advertisement Commands

- ```

config
 — service
 — vprn
 — [no] router-advertisement
 — [no] dns-options
 — server ipv6-address
 — no server
 — rdnss-lifetime {seconds | infinite}
 — no rdnss-lifetime
 — [no] interface ip-int-name
 — current-hop-limit number
 — no current-hop-limit
 — [no] dns-options
 — server ipv6-address
 — no server
 — rdnss-lifetime seconds
 — no rdnss-lifetime
 — [no] include-dns
 — [no] managed-configuration
 — max-advertisement-interval seconds
 — no max-advertisement-interval
 — min-advertisement-interval seconds
 — no min-advertisement-interval
 — mtu mtu-bytes
 — no mtu
 — [no] other-stateful-configuration
 — prefix [ipv6-prefix/prefix-length]
 — no prefix
 — [no] autonomous
 — [no] on-link
 — preferred-lifetime {seconds | infinite}
 — no preferred-lifetime
 — valid-lifetime {seconds | infinite}
 — no valid-lifetime
 — reachable-time milli-seconds
 — no reachable-time
 — retransmit-time milli-seconds
 — no retransmit-time
 — router-lifetime seconds
 — no router-lifetime
 — [no] shutdown
 — [no] use-virtual-mac

```

### 3.8.1.13 NTP Commands

The **ntp-server** command is not supported in the **vprn ntp** context. Then NTP is configured in a VPRN service, the NTP server mode is assumed and is not optional.

```

config
 — service
 — vprn
 — [no] ntp
 — [no] authenticate
 — [no] authentication-check
 — authentication-key key-id key key [hash | hash2] type {des |
 message-digest}
 — no authentication-key key-id
 — [no] broadcast [router router-name] {interface ip-int-name} [key-id
 key-id] [version version] [ttl ttl]

```

### 3.8.1.14 NAT Commands

```

config
 — service
 — vprn service-id [customer customer-id]
 — no vprn service-id
 — [no] nat
 — inside
 — [no] destination-prefix ip-prefix/length
 — dual-stack-lite
 — [no] address ipv6-address
 — tunnel-mtu mtu-bytes
 — no tunnel-mtu
 — [no] shutdown
 — subscriber-prefix-length prefix-length
 — no subscriber-prefix-length
 — l2-aware
 — [no] address ip-address/mask
 — nat-policy nat-policy-name
 — no nat-policy
 — redundancy
 — peer ip-address
 — no peer
 — steering-route ip-prefix/length
 — no steering-route
 — outside
 — pool nat-pool-name [nat-group nat-group-id type pool-type [no-
 allocate] [create]
 — no pool nat-pool-name
 — address-range start-ip-addr end-ip-addr [create]
 — no address-range start-ip-address end-ip-address
 — description description-string
 — no description

```

- [no] **drain**
- **description** *description-string*
- **no description**
- **mode** {auto | n apt}
- **no mode**
- **port-forwarding-range** *range-end*
- **no port-forwarding-range**
- **port-reservation blocks** *num-blocks*
- **port-reservation ports** *num-ports*
- **no port-reservation**
- **redundancy**
  - **export** *ip-prefix/length*
  - **no export**
  - **follow router** *router-instance pool name*
  - **no follow**
  - **monitor** *ip-prefix/length*
  - **no monitor**
- [no] **shutdown**
- **subscriber-limit** *limit*
- **no subscriber-limit**
- **watermarks high** *percentage-high low percentage-low*
- **no watermarks**

### 3.8.1.15 Subscriber Interface Commands

- ```

config
  — service
    — vpri service-id [customer customer-id]
    — no vpri service-id
    — subscriber-interface ip-int-name [fwd-service service-id fwd-subscriber-
      interface ip-int-name] [create]
    — no subscriber-interface ip-int-name
      — [no] address {ip-address/mask | ip-address netmask} [gw-ip-address
        ip-address] [populate-host-routes] [track-srrp srrp-instance
        [holdup-time msec]]
      — [no] allow-unmatching-subnets
      — authentication-policy name
      — no authentication-policy
      — description long-description-string
      — no description
      — dhcp
        — client-applications dhcp ppp
        — no client-applications
        — description description-string
        — no description
        — gi-address ip-address [src-ip-addr]
        — no gi-address
        — lease-populate nbr-of-leases
        — no lease-populate
        — [no] option
          — [no] vendor-specific-option

```

```

— [no] client-mac-address
— [no] sap-id
— [no] service-id
— string text
— no string
— [no] system-id
— proxy-server
  — emulated-server ip-address
  — no emulated-server
  — lease-time [days days] [hrs hours] [min minutes] [sec
    seconds] [radius-override]
  — no lease-time
  — [no] shutdown
— python-policy name
— no python-policy
— relay-proxy [release-update-src-ip] [siaddr-override ip-
  address]
— no relay-proxy
— server server1 [server2]
— no server
— [no] shutdown
— [no] group-interface ip-int-name
  — arp-host
    — host-limit max-num-hosts
    — no host-limit
    — min-auth-interval min-auth-interval
    — no min-auth-interval
    — sap-host-limit max-num-hosts-sap
    — no sap-host-limit
    — [no] shutdown
  — [no] arp-populate
  — arp-timeout seconds
  — no arp-timeout
  — authentication-policy name
  — no authentication-policy
  — description long-description-string
  — no description
  — dhcp
    — client-applications dhcp ppp
    — no client-applications
    — description description-string
    — no description
    — filter filter-id
    — no filter
    — gi-address ip-address [src-ip-addr]
    — no gi-address
    — lease-populate nbr-of-leases
    — no lease-populate
    — [no] match-circuit-id
    — [no] option
      — action {replace | drop | keep}
      — no action
      — circuit-id [ascii-tuple | ifindex | sap-id | vlan-
        ascii-tup]

```

-
- **no circuit-id**
 - **remote-id** [mac | string *string*]
 - **no remote-id**
 - [no] **vendor-specific-option**
 - [no] **client-mac-address**
 - [no] **sap-id**
 - [no] **service-id**
 - **string** text
 - **no string**
 - [no] **system-id**
 - **proxy-server**
 - **emulated-server** ip-address
 - **no emulated-server**
 - **lease-time** [days *days*] [hrs *hours*] [min *minutes*] [sec *seconds*] [override]
 - **no lease-time**
 - [no] **shutdown**
 - **python-policy** *name*
 - **no python-policy**
 - **relay-proxy** [release-update-src-ip] [siaddr-override ip-address]
 - **no relay-proxy**
 - **server** server1 [server2]
 - **no server**
 - [no] **shutdown**
 - [no] **trusted**
 - **user-db** local-user-db-name
 - **no user-db**
 - [no] **enable-ingress-stats**
 - **host-connectivity-verify** [interval *interval*] [action {remove | alarm}] [timeout *retry-timeout*] [retry-count *count*] [family *family*]
 - **hold-time**
 - **up** ip seconds
 - **no up** ip
 - **up** ipv6 seconds
 - **no up** ipv6
 - **down** ip seconds [init-only]
 - **no down** ip
 - **down** ipv6 seconds [init-only]
 - **no down** ipv6
 - **icmp**
 - [no] **mask-reply**
 - **redirects** [number seconds]
 - **no redirects**
 - **ttl-expired** [number seconds]
 - **no ttl-expired**
 - **unreachables** [number seconds]
 - **no unreachables**
 - [no] **ipv6**
 - [no] **allow-unmatching-prefixes**
 - **delegated-prefix-length** bits
 - **delegated-prefix-length** variable
 - **no delegated-prefix-length**

-
- [no] **dhcp6**
 - [no] **option**
 - **interface-id**
 - **interface-id** **ascii-tuple**
 - **interface-id** **ifindex**
 - **interface-id** **sap-id**
 - **interface-id** **string**
 - **no interface-id**
 - [no] **remote-id**
 - [no] **router-advertisements**
 - **no current-hop-limit**
 - [no] **managed-configuration**
 - **max-advertisement** *seconds*
 - **no max-advertisement**
 - **min-advertisement** *seconds*
 - **no min-advertisement**
 - **mtu** *bytes*
 - **no mtu**
 - [no] **other-stateful-configuration**
 - [no] **prefix-options**
 - [no] **autonomous**
 - **preferred-lifetime** [*seconds* | **infinite**]
 - **no preferred-lifetime**
 - **valid-lifetime** [*seconds* | **infinite**]
 - **no valid-lifetime**
 - **reachable-time** *milliseconds*
 - **no reachable-time**
 - **retransmit-time** *milliseconds*
 - **no retransmit-time**
 - **router-lifetime** *seconds*
 - **router-lifetime** **no-default-router**
 - **no router-lifetime**
 - [no] **proxy-server**
 - **renew-timer** *seconds*
 - **no renew-timer**
 - **rebind-timer** *seconds*
 - **no rebind-timer**
 - **preferred-lifetime** [*seconds* | **infinite**]
 - **no preferred-lifetime**
 - **valid-lifetime** [*seconds* | **infinite**]
 - **no valid-lifetime**
 - **client-applications** [**dhcp**] [**ppp**]
 - **no client-applications**
 - [no] **vpn**
 - [no] **router-advertisements**
 - **current-hop-limit** *hop-count*
 - [no] **local-proxy-arp**
 - [no] **mac** *ieee-address*
 - [no] **pppoe**
 - **description** *description-string*
 - **no description**
 - **dhcp-client**
 - [no] **ccag-use-origin-sap**
 - **pap-chap-user-db** *local-user-db-name*

```

— no pap-chap-user-db
— policy pppoe-policy-name
— no policy
— sap-session-limit sap-session-limit
— no sap-session-limit
— session-limit session-limit
— no session-limit
— user-db local-user-db-name
— no user-db
— [no] shutdown
— [no] proxy-arp-policy policy-name [policy-name]
— broadcast red-ip-int-name
— no redundant-interface
— [no] remote-proxy-arp
— [no] sap sap-id
    — accounting-policy acct-policy-id
    — no accounting-policy [acct-policy-id]
    — anti-spoof {ip | ip-mac | nh-mac}
    — no anti-spoof
    — app-profile app-profile-name
    — no app-profile
    — atm
        — egress
            — traffic-desc traffic-desc-profile-id
            — no traffic-desc
        — encapsulation atm-encap-type
        — ingress
            — traffic-desc traffic-desc-profile-id
            — no traffic-desc
        — oam
            — [no] alarm-cells
            — [no] periodic-loopback
    — calling-station-id calling-station-id
    — no calling-station-id
    — [no] bfd-enable
    — cpu-protection {[mac-monitoring] | [eth-cfm-
        monitoring [aggregate] [car]]}
    — no cpu-protection
    — default-host ip-address/mask next-hop next-hop-ip
    — no default-host ip-address/mask
    — description long-description-string
    — no description
    — dist-cpu-protection policy-name
    — no dist-cpu-protection
    — egress
        — [no] agg-rate
            — [no] limit-unused-bandwidth
            — [no] queue-frame-based-accounting
            — rate {max | rate}
            — no rate
        — filter ip ip-filter-id
        — filter ipv6 ipv6-filter-id
        — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
        — no filter

```

-
- **policer-control-policy** *policy-name*
 - **no policer-control-policy**
 - **[no] qinq-mark-top-only**
 - **qos** *policy-id* [**port-redirect-group** *queue-group-name* **instance** *instance-id*]
 - **no qos**
 - **scheduler-policy** *scheduler-policy-name*
 - **no scheduler-policy**
 - **[no] fwd-wholesale**
 - **pppoe** *service-id*
 - **no pppoe**
 - **igmp-host-tracking**
 - **expiry-time** *expiry-time*
 - **no expiry-time**
 - **import** *policy-name*
 - **no import**
 - **max-num-groups** *max-num-groups*
 - **no max-num-groups**
 - **max-num-grp-sources** [*number*]
 - **no max-num-grp-sources**
 - **max-num-sources** *max-num-sources*
 - **no max-num-sources**
 - **ingress**
 - **filter ip** *ip-filter-id*
 - **filter ipv6** *ipv6-filter-id*
 - **no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
 - **no filter**
 - **match-qinq-dot1p** {**top** | **bottom**}
 - **no match-qinq-dot1p**
 - **qos** *policy-id* [**shared-queuing** | **multipoint-shared**] [**fp-redirect-group** *queue-group-name* **instance** *instance-id*]
 - **no qos**
 - **scheduler-policy** *scheduler-policy-name*
 - **no scheduler-policy**
 - **lag-link-map-profile** *link-map-profile-id*
 - **no lag-link-map-profile**
 - **multi-service-site** *customer-site-name*
 - **no multi-service-site**
 - **static-host ip** *ip/did-address* [**mac** *ieee-address*] [**create**]
 - **static-host mac** *ieee-address* [**create**]
 - **no static-host** [**ip** *ip-address*] **mac** *ieee-address*
 - **no static-host all** [**force**]
 - **no static-host ip** *ip-address*
 - **ancp-string** *ancp-string*
 - **no ancp-string**
 - **app-profile** *app-profile-name* [**scope** *scope-type*]
 - **no app-profile**
 - **inter-dest-id** *intermediate-destination-id*
 - **no inter-dest-id**
 - **managed-routes**

```

— route {ip-prefix/length | ip-prefix netmask}
  [create]
— no route {ip-prefix/length | ip-prefix
  netmask}
— [no] shutdown
— sla-profile sla-profile-name
— no sla-profile
— sub-profile sub-profile-name
— no sub-profile
— subscriber sub-ident
— no subscriber
— [no] subscriber-sap-id
— [no] shutdown
— [no] sub-sla-mgmt
  — def-sla-profile default-sla-profile-name
  — no def-sla-profile
  — def-sub-profile default-subscriber-profile-name
  — no def-sub-profile
  — multi-sub-sap subscriber-limit
  — no multi-sub-sap
  — [no] shutdown
  — single-sub-parameters
    — non-sub-traffic sub-profile sub-profile-
      name sla-profile sla-profile-name
      [subscriber sub-ident-string]
    — no non-sub-traffic
    — [no] profiled-traffic-only
  — sub-ident-policy sub-ident-policy-name
  — no sub-ident-policy
— shcv-policy-ipv4 policy-name
— no shcv-policy-ipv4
— shcv-policy-ipv6 policy-name
— no shcv-policy-ipv6
— [no] shutdown
— [no] srrp srrp-id
— [no] bfd-enable svc-id interface interface-name dst-ip
  ip-address
— [no] bfd-enable interface interface-name dst-ip ip-
  address
— [no] bfd-enable interface interface-name dst-ip ip-
  address name name
— description description-string
— no description
— gw-mac mac-address
— no gw-mac
— keep-alive-interval interval
— no keep-alive-interval
— message-path sap-id
— no message-path
— [no] policy vrrp-policy-id
— priority priority
— no priority
— send-fib-population-packets all | outer-tag-only
— no send-fib-population-packets

```

- [no] **shutdown**
- [no] **wpp**
 - **initial-app-profile** *profile-name*
 - **no initial-app-profile**
 - **initial-sla-profile** *profile-name*
 - **no initial-sla-profile**
 - **initial-sub-profile** *profile-name*
 - **no initial-sub-profile**
 - **portal router** *router-instance name wpp-portal-name*
 - **no portal**
 - [no] **restore-disconnected**
 - [no] **shutdown**
- [no] **urpf-check**
 - **mode** {strict | loose | strict-no-ecmp}
- [no] **ipv6**
 - [no] **delegated-prefix-length** *prefix-length*
 - [no] **subscriber-prefixes**
 - **prefix** *ipv6-address/prefix-length [pd] [wan-host]*
 - **no prefix** *ipv6-address/prefix-length*
- [no] **private-retail-subnets**
- [no] **shutdown**

3.8.1.15.1 Group Interface SAP ETH-CFM Commands

- ```
config>service>vprn>sub-if>grp-if>sap
```
- **eth-cfm**
    - [no] **collect-lmm-stats**
    - **collect-lmm-fc-stats**
      - **fc** *fc-name [fc-name]*
      - **no fc**
      - **fc-in-profile** *fc-name [fc-name]*
      - **no fc-in-profile**
    - **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**
      - [no] **ccm-enable**
      - **ccm-ltm-priority** *priority*
      - **no ccm-ltm-priority**
      - [no] **description**
      - [no] **eth-test-enable**
      - [no] **test-pattern** {all-zeros | all-ones} [crc-enable]
      - **fault-propagation-enable** {use-if-tlv | suspend-ccm}
      - **no fault-propagation-enable**
      - **grace**
        - **eth-ed**
          - **max-rx-defect-window** *seconds*
          - **no max-rx-defect-window**
          - **priority** *priority*
          - **no priority**
          - [no] **rx-eth-ed**
          - [no] **tx-eth-ed**

- **eth-vsm-grace**
  - [no] **rx-eth-vsm-grace**
  - [no] **tx-eth-vsm-grace**
- **low-priority-defect** {allDef | macRemErrXcon | remErrXcon | errXcon | xcon | noXcon}
- **one-way-delay-threshold** *seconds*
- [no] **shutdown**
- **squelch-ingress-levels** [md-level [*md-level...*]]
- **no squelch-ingress-levels**
- **tunnel-fault** [accept | ignore]

### 3.8.1.16 Interface Commands

- ```

config
  — service
    — vprn
      — [no] interface ip-int-name
        — [no] active-cpm-protocols
        — address [ip-address/mask | ip-address netmask] [broadcast all-ones | host-ones] [track-srrp srrp-instance]
        — no address [ip-address/mask | ip-address netmask]
        — [no] allow-directed-broadcasts
        — [no] arp-learn-unsolicited
        — arp-limit limit [log-only] [threshold percent]
        — no arp-limit
        — [no] arp-populate
        — [no] arp-populate-host-route
        — [no] arp-proactive-refresh
        — arp-retry-timer timer-multiple
        — no arp-retry-timer
        — arp-route-tag tag
        — no arp-route-tag
        — arp-timeout [seconds]
        — no arp-timeout
        — authentication-policy name
        — no authentication-policy
        — bfd transmit-interval [receive receive-interval] [multiplier multiplier] [echo-receive echo-interval] [type cpm-np]
        — no bfd
        — cflowd-parameters
        — no cflowd-parameters
          — sampling {unicast | multicast} type {acl | interface} [direction {ingress-only | egress-only | both}]
          — no sampling {unicast | multicast}
        — cpu-protection policy-id
        — no cpu-protection
        — description long-description-string
        — no description
        — dhcp
          — description description-string
          — no description

```

-
- **gi-address** *ip-address* [*src-ip-addr*]
 - **no gi-address**
 - **lease-populate** [*nbr-of-leases*]
 - **no lease-populate**
 - **[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**
 - **proxy-server**
 - **emulated-server** *ip-address*
 - **no emulated-server**
 - **lease-time** [*days days*] [*hrs hours*] [*min minutes*] [*sec seconds*] [*radius-override*]
 - **no lease-time**
 - **[no] shutdown**
 - **[no] relay-plain-bootp**
 - **relay-proxy** [*release-update-src-ip*] [*siaddr-override ip-address*]
 - **no relay-proxy**
 - **python-policy** *policy-name*
 - **no python-policy**
 - **sap** *sap-id*
 - **server** *server1* [*server2*]
 - **no server**
 - **[no] shutdown**
 - **[no] trusted**
 - **[no] use-arp**
 - **dynamic-tunnel-redundant-next-hop** *ip-address*
 - **no dynamic-tunnel-redundant-next-hop**
 - **[no] enable-ingress-stats**
 - **[no] enable-mac-accounting**
 - **host-connectivity-verify** [*source {vrrp | interface}*] [*interval interval*] [*action {remove | alarm}*] [*timeout retry-timeout*] [*retry-count count*]
 - **hold-time**
 - **up** *ip seconds*
 - **no up** *ip*
 - **up** *ipv6 seconds*
 - **no up** *ipv6*
 - **down** *ip seconds* [*init-only*]
 - **no down** *ip*
 - **down** *ipv6 seconds* [*init-only*]
 - **no down** *ipv6*

-
- **icmp**
 - [no] **mask-reply**
 - **param-problem** *number seconds*
 - no **param-problem** [*number seconds*]
 - **redirects** *number seconds*
 - no **redirects** [*number seconds*]
 - **ttl-expired** *number seconds*
 - no **ttl-expired** [*number seconds*]
 - **unreachables** *number seconds*
 - no **unreachables** [*number seconds*]
 - **if-attribute**
 - [no] **admin-group** *group-name* [*group-name*]
 - **srfg-group** *group-name* [*group-name*]
 - no **srfg-group**
 - **ingress**
 - **policy-accounting** *template-name*
 - no **policy-accounting**
 - **ip-helper-address** *gateway-address*
 - no **ip-helper-address**
 - **ip-mtu** *octets*
 - no **ip-mtu**
 - **ipcp**
 - **dns** *ip-address* [**secondary** *ip-address*]
 - **dns secondary** *ip-address*
 - no **dns** [*ip-address*] [**secondary** *ip-address*]
 - **peer-ip-address** *ip-address*
 - no **peer-ip-address**
 - [no] **ipv6**
 - **address** *ipv6-address/prefix-length* [**eui-64**] [**track-srrp** *srrp-instance*] [**modifier** *cga-modifier*] [**dad-disable**] [**primary-preference** *primary-preference*]
 - no **address** *ipv6-address/prefix-length*
 - **bfd** *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier*] [**echo-receive** *echo-interval*] [**type** *cpm-np*]
 - no **bfd**
 - [no] **dad-disable**
 - [no] **dhcp6-relay**
 - **lease-populate** [*nbr-of-leases*]
 - no **lease-populate**
 - [no] **neighbor-resolution**
 - [no] **dhcp6-server**
 - **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**
 - **unreachables** *number seconds*
 - no **unreachables**
 - **link-local-address** *ipv6-address* [**dad-disable**]
 - no **link-local-address**

-
- [no] **local-proxy-nd**
 - **nd-learn-unsolicited** {global | link-local | both}
 - no **nd-learn-unsolicited**
 - [no] **nd-populate-host-route**
 - **nd-proactive-refresh** {global | link-local | both}
 - no **nd-proactive-refresh**
 - **nd-route-tag** tag
 - no **nd-route-tag**
 - **neighbor** ipv6-address mac-address
 - no **neighbor** ipv6-address
 - **neighbor-limit** limit [log-only] [threshold percent]
 - no **neighbor-limit**
 - **proxy-nd-policy** policy-name [policy-name]
 - no **proxy-nd-policy**
 - **python-policy** policy-name
 - no **python-policy**
 - [no] **qos-route-lookup**
 - **stale-time** seconds
 - no **stale-time**
 - [no] **secure-nd**
 - [no] **allow-unsecured-msgs**
 - **link-local-modifier** modifier
 - no **link-local-modifier**
 - **public-key-min-bits** bits
 - no **public-key-min-bits**
 - **security-parameter** sec
 - no **security-parameter**
 - [no] **shutdown**
 - **tcp-mss** mss-value
 - [no] **tcp-mss**
 - [no] **urpf-check**
 - **mode** {strict | loose | strict-no-ecmp}
 - no **mode**
 - **load-balancing**
 - **egr-ip-load-balancing** {source | destination | inner-ip}
 - no **egr-ip-load-balancing**
 - [no] **spi-load-balancing**
 - [no] **teid-load-balancing**
 - **local-dhcp-server** local-server-name
 - no **local-dhcp-server**
 - [no] **local-proxy-arp**
 - [no] **loopback**
 - **mac** ieee-address
 - no **mac** [ieee-address]
 - **monitor-oper-group** name
 - no **monitor-oper-group**
 - [no] **proxy-arp-policy**
 - [no] **ptp-hw-assist**
 - **qos-route-lookup** [source | destination]
 - no **qos-route-lookup**
 - [no] **remote-proxy-arp**
 - **secondary** {ip-address/mask | ip-address netmask} [broadcast all-ones | host-ones] [igp-inhibit]
 - no **secondary** {ip-address/mask | ip-address netmask}

- **shcv-policy-ipv4** *policy-name*
- **no shcv-policy-ipv4**
- **[no] shutdown**
- **static-arp** *ip-address* *ieee-address*
- **no static-arp** *ip-address* [*ieee-address*]
- **static-tunnel-redundant-next-hop** *ip-address*
- **no static-tunnel-redundant-next-hop**
- **tcp-mss** *mss-value*
- **[no] tcp-mss**
- **tos-marking-state** {*trusted* | *untrusted*}
- **no tos-marking-state**
- **unnumbered** [*ip-int-name* | *ip-address*]
- **no unnumbered**
- **[no] urpf-check**
 - **mode** {*strict* | *loose* | *strict-no-ecmp*}
 - **no mode**
- **vas-if-type** {*to-from-access* | *to-from-network* | *to-from-both*}
- **no vas-if-type**
- **vpls** *service-name*
- **no vpls**
 - **egress**
 - **reclassify-using-qos** *policy-id*
 - **no reclassify-using-qos**
 - **v4-routed-override-filter** *ipv4-filter-id*
 - **no v4-routed-override-filter**
 - **v6-routed-override-filter** *ipv6-filter-id*
 - **no v6-routed-override-filter**
 - **ingress**
 - **v4-routed-override-filter** *ipv4-filter-id*
 - **no v4-routed-override-filter**
 - **v6-routed-override-filter** *ipv6-filter-id*
 - **no v6-routed-override-filter**

3.8.1.17 Network Interface Commands

- ```

config
 — service
 — vprn
 — network-interface interface-name [create]
 — no network-interface interface-name
 — address ip-address [/mask] [/netmask] [broadcast {all-ones | host-ones}]
 — no address
 — [no] allow-directed-broadcasts
 — [no] arp-populate
 — arp-retry-timer timer-multiple
 — no arp-retry-timer
 — arp-timeout [seconds]
 — no arp-timeout
 — bfd transmit-interval [receive receive-interval] [multiplier multiplier]
 [echo-receive echo-interval] [type cpm-np]

```

- 
- **no bfd**
  - **cflowd-parameters**
  - **no cflowd-parameters**
    - **sampling** {unicast | multicast} type {acl | interface} [direction {ingress-only | egress-only | both}]
    - **no sampling** {unicast | multicast}
  - **cpu-protection** *policy-id* [mac-monitoring] | [eth-cfm-monitoring [aggregate] [car]]
  - **no cpu-protection**
  - **description** *long-description-string*
  - **no description**
  - **dist-cpu-protection** *policy-name*
  - **no dist-cpu-protection**
  - **egress**
    - **filter** ip *ip-filter-id*
    - **filter** ipv6 *ipv6-filter-id*
    - **no filter** [ip *ip-filter-id*] [ipv6 *ipv6-filter-id*]
  - **hold-time**
    - **up** ip *seconds*
    - **no up** ip
    - **up** ipv6 *seconds*
    - **no up** ipv6
    - **down** ip *seconds* [init-only]
    - **no down** ip
    - **down** ipv6 *seconds* [init-only]
    - **no down** ipv6
  - **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*]
  - **ingress**
    - **filter** ip *ip-filter-id*
    - **filter** ipv6 *ipv6-filter-id*
    - **no filter** [ip *ip-filter-id*] [ipv6 *ipv6-filter-id*]
  - **ip-mtu** *octets*
  - **no ip-mtu**
  - **lag** *lag-id[:encap-val]*
  - **no lag**
  - **lag-per-link-hash** class {1 | 2 | 3} weight *weight*
  - **no lag-per-link-hash**
  - [no] **loopback**
  - **no load-balancing**
    - **egr-ip-load-balancing** {source | destination | inner-ip}
    - **no egr-ip-load-balancing**
    - **lsr-load-balancing** *hashing-algorithm*
    - **no lsr-load-balancing**
    - [no] **sbi-load-balancing**
    - [no] **teid-load-balancing**
  - **mac** *ieee-address*
  - **no mac**



- **qos** *network-policy-id* **port-redirect-group** *queue-group-name* *egress-instance instance-id* *fp-redirect-group* *queue-group-name* *ingress-instance instance-id*
- **no qos**
- **secondary** {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**] [**igmp-inhibit**]
- **no secondary** {*ip-address/mask* | *ip-address netmask*}
- **static-arp** *ieee-mac-address* *unnumbered*
- **no static-arp** *unnumbered*
- **tos-marking-state** {**trusted** | **untrusted**}
- **no tos-marking-state**
- [no] **urpf-check**
  - **mode** {**strict** | **loose** | **strict-no-ecmp**}

### 3.8.1.18 Interface Spoke SDP Commands

- ```

config
  — service
    — vprn service-id [customer customer-id]
    — no vprn service-id
    — [no] interface ip-int-name
      — spoke-sdp sdp-id [:vc-id] vc-type {ether | ipipe} [create]
      — no spoke-sdp sdp-id [:vc-id] vc-type {ether | ipipe} [create]
        — aarp aarpId type type
        — no aarp
        — accounting-policy acct-policy-id
        — no accounting-policy
        — app-profile app-profile-name
        — no app-profile
        — bfd-template name
        — no bfd-template
        — [no] bfd-enable
        — [no] control-channel-status
          — [no] acknowledgment
          — refresh-timer value
          — no refresh-timer
          — request-timer timer1 retry-timer timer2 [timeout-multiplier multiplier]
          — no request-timer
        — [no] control-word
        — egress
          — filter ip ip-filter-id
          — filter ipv6 ipv6-filter-id
          — no filter
          — qos network-policy-id port-redirect-group queue-group-name [instance instance-id]
          — no qos
          — vc-label egress-vc-label
          — no vc-label [egress-vc-label]
        — [no] entropy-label
        — eth-cfm

```

-
- [no] **collect-lmm-stats**
 - **collect-lmm-fc-stats**
 - **fc** *fc-name* [*fc-name*]
 - **no fc**
 - **fc-in-profile** *fc-name* [*fc-name*]
 - **no fc-in-profile**
 - **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**
 - [no] **interface-support-enable**
 - [no] **ccm-enable**
 - **ccm-ltm-priority** *priority*
 - **no ccm-ltm-priority**
 - **grace**
 - **eth-ed**
 - **max-rx-defect-window** *seconds*
 - **no max-rx-defect-window**
 - **priority** *priority*
 - **no priority**
 - [no] **rx-eth-ed**
 - [no] **tx-eth-ed**
 - **eth-vsm-grace**
 - [no] **rx-eth-vsm-grace**
 - [no] **tx-eth-vsm-grace**
 - **ccm-padding-size** *ccm-padding*
 - **no ccm-padding-size** *ccm-padding*
 - [no] **csf-enable**
 - **multiplier** *multiplier-value*
 - **no multiplier**
 - **description** *description-string*
 - **no description**
 - [no] **eth-test-enable**
 - [no] **test-pattern** {*all-zeros* | *all-ones*} [*crc-enable*]
 - **fault-propagation-enable** {*use-if-tilv* | *suspend-ccm*}
 - **no fault-propagation-enable**
 - **low-priority-defect** {*allDef* | *macRemErrXcon* | *remErrXcon* | *errXcon* | *xcon* | *noXcon*}
 - **one-way-delay-threshold** *seconds*
 - [no] **snqlch-ingress-levels** [*md-level* [*md-level*...]]
 - [no] **hash-label**
 - **qos**
 - **filter ip** *ip-filter-id*
 - **filter ipv6** *ipv6-filter-id*}
 - **no filter**
 - **ingress** *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*
 - **no qos**
 - **vc-label** *ingress-vc-label*
 - **no vc-label** [*ingress-vc-label*]
 - [no] **shutdown**

- **transit-policy** {*ip ip-aasub-policy-id* | *prefix prefix-aasub-policy-id*}
- **no transit-policy**
- [no] **pw-path-id**
 - **agi** *agi*
 - **no agi**
 - **saii-type2** *global-id:node-id:ac-id*
 - **no saii-type2**
 - **taii-type2** *global-id:node-id:ac-id*
 - **no taii-type2**

3.8.1.19 Interface VRRP Commands

- ```

config
 — service
 — vrpn service-id [customer customer-id]
 — no vrpn service-id
 — interface ip-int-name
 — vrrp
 — vrrp virtual-router-id [owner] [passive]
 — no vrrp virtual-router-id
 — [no] backup ip-address
 — [no] bfd-enable interface interface-name dst-ip ip-address
 — [no] bfd-enable interface interface-name dst-ip ip-address name name
 — [no] bfd-enable svc-id interface interface-name dst-ip ip-address
 — init-delay seconds
 — no init-delay
 — mac ieee-address
 — no mac
 — [no] master-int-inherit
 — message-interval {[seconds] [milliseconds]}
 — 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
 — ipv6
 — vrrp virtual-router-id [owner] [passive]
 — no vrrp virtual-router-id

```

- **authentication-key** {*authentication-key* | *hash-key*}  
[*hash* | *hash2*]
- **no authentication-key**
- [no] **backup** *ip-address*
- [no] **bfd-enable** *svc-id* **interface** *interface-name* **dst-ip** *ip-address*
- [no] **bfd-enable** **interface** *interface-name* **dst-ip** *ip-address*
- [no] **bfd-enable** **interface** *interface-name* **dst-ip** *ip-address* **name** *name*
- **init-delay** *seconds*
- **no init-delay**
- **mac** *ieee-address*
- **no mac**
- [no] **master-int-inherit**
- **message-interval** {[*seconds*] [*milliseconds*]  
*milliseconds*}
- **no message-interval**
- **oper-group** *group-name*
- **no oper-group**
- [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**

### 3.8.1.20 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
          — aarp aarpId type type
          — no aarp
          — accounting-policy acct-policy-id
          — no accounting-policy [acct-policy-id]
          — anti-spoof {ip | mac | ip-mac}
          — no anti-spoof
          — app-profile app-profile-name
          — no app-profile
          — atm
            — egress
              — traffic-desc traffic-desc-profile-id

```

```

    — no traffic-desc
— encapsulation atm-encap-type
— ingress
    — traffic-desc traffic-desc-profile-id
    — no traffic-desc
— oam
    — [no] alarm-cells
    — [no] periodic-loopback
— bandwidth bandwidth
— no bandwidth
— calling-station-id calling-station-id
— no calling-station-id
— [no] bfd-enable
— cpu-protection policy-id [mac-monitoring] | [eth-cfm-
    monitoring [aggregate] [car]]
— no cpu-protection
— description long-description-string
— no description
— dist-cpu-protection policy-name
— no dist-cpu-protection
— egress
    — [no] agg-rate
        — [no] limit-unused-bandwidth
        — [no] queue-frame-based-accounting
        — rate {max | rate}
        — no rate
    — filter ip ip-filter-id
    — no filter [ip ip-filter-id]
    — [no] hsmdda-queue-override
        — secondary-shaper secondary-shaper-name
        — no secondary-shaper
        — wrr-policy hsmdda-wrr-policy-name
        — no wrr-policy
        — packet-byte-offset {add add-bytes | subtract
            sub-bytes}
        — no packet-byte-offset
        — queue queue-id
        — no queue queue-id
            — wrr-weight weight
            — no wrr-weight
            — mbs size {[bytes | kilobytes] | default}
            — no mbs
            — [no] monitor-depth
            — rate pir-rate
            — no rate
            — slope-policy hsmdda-slope-policy-name
                allowable
            — no slope-policy
    — [no] policer-override
        — policer policer-id [create]
        — no policer policer-id
— source ip-address
— remote-ip ip-address
— backup-remote-ip ip-address

```

-
- [no] **qinq-mark-top-only**
 - **qos** *policy-id* [**port-redirect-group** *queue-group-name*
instance *instance-id*]
 - no **qos** [*policy-id*]
 - **queue-group-redirect-list** *redirect-list-name*
 - no **queue-group-redirect-list**
 - [no] **queue-override**
 - **hs-secondary-shaper** *policy-name*
 - no **hs-secondary-shaper**
 - **hs-wrr-group** *group-id* [**create**]
 - no **hs-wrr-group**
 - **class-weight** *weight*
 - no **class-weight**
 - **percent-rate** *percent*
 - no **percent-rate**
 - **rate** *rate*
 - no **rate**
 - [no] **queue** *queue-id*
 - **adaptation-rule** [**pir** *adaptation-rule*] [**cir**
adaptation-rule]
 - no **adaptation-rule**
 - **avg-frame-overhead** *percentage*
 - no **avg-frame-overhead**
 - **burst-limit** {**default** | *size* [**bytes** |
kilobytes]}
 - no **burst-limit**
 - **cbs** *size-in-kbytes*
 - no **cbs**
 - **drop-tail** *low*
 - **percent-reduction-from-mbs**
percent
 - no **percent-reduction-from-mbs**
 - **hs-class-weight** *weight*
 - no **hs-class-weight**
 - **hs-wred-queue** *policy* *slope-policy-name*
 - no **hs-wred-queue**
 - **hs-wrr-weight** *weight*
 - no **hs-wrr-weight**
 - **mbs** *size* {[**bytes** | **kilobytes**] | **default**}
 - no **mbs**
 - **parent** [*weight* *weight*] [**cir-weight** *cir-weight*]
 - no **parent**
 - **percent-rate** *pir-percent* [**cir** *cir-percent*]
 - no **percent-rate**
 - **rate** *pir-rate* [**cir** *cir-rate*]
 - no **rate**
 - [no] **scheduler-override**
 - [no] **scheduler** *scheduler-name*
 - **parent** [*weight* *weight*] [**cir-weight** *cir-weight*]
 - no **parent**
 - **rate** *pir-rate* [**cir** *cir-rate*]
 - no **rate**

-
- **scheduler-policy** *scheduler-policy-name*
 - **no scheduler-policy**
 - **eth-cfm**
 - [no] **collect-lmm-stats**
 - **collect-lmm-fc-stats**
 - **fc** *fc-name* [*fc-name*]
 - **no fc**
 - **fc-in-profile** *fc-name* [*fc-name*]
 - **no fc-in-profile**
 - **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**
 - [no] **interface-support-enable**
 - [no] **ccm-enable**
 - **ccm-ltm-priority** *priority*
 - **no ccm-ltm-priority**
 - [no] **ccm-padding-size** *ccm-padding*
 - [no] **csf-enable**
 - **multiplier** *multiplier-value*
 - **no multiplier**
 - **description** *description-string*
 - **no description**
 - [no] **eth-test-enable**
 - [no] **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
 - **fault-propagation-enable** {**use-if-tlv** | **suspend-ccm**}
 - **no fault-propagation-enable**
 - **grace**
 - **eth-ed**
 - **max-rx-defect-window** *seconds*
 - **no max-rx-defect-window**
 - **priority** *priority*
 - **no priority**
 - [no] **rx-eth-ed**
 - [no] **tx-eth-ed**
 - **eth-vsm-grace**
 - [no] **rx-eth-vsm-grace**
 - [no] **tx-eth-vsm-grace**
 - **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}
 - **one-way-delay-threshold** *seconds*
 - **snelch-ingress-levels** [*md-level* [*md-level...*]]
 - **no snelch-ingress-levels**
 - **tunnel-fault** [**accept** | **ignore**]
 - **frame-relay**
 - [no] **frf-12**
 - **ete-fragment-threshold** *threshold*
 - **no ete-fragment-threshold**
 - [no] **interleave**
 - **scheduling-class** *class-id*
 - **no scheduling-class**
 - [no] **scheduling-class**

-
- **host-lockout-policy** *policy-name*
 - **no host-lockout-policy**
 - **[no] host-shutdown**
 - **ingress**
 - **filter** *ip ip-filter-id*
 - **no filter** [*ip ip-filter-id*]
 - **match-qinq-dot1p** {*top* | *bottom*}
 - **[no] policer-override**
 - **policer** *policer-id* [*create*]
 - **no policer** *policer-id*
 - **qos** *policy-id* [*shared-queuing* | *multipoint-shared*] [*fp-redirect-group queue-group-name instance instance-id*]
 - **no qos** [*policy-id*]
 - **queue-group-redirect-list** *redirect-list-name*
 - **no queue-group-redirect-list**
 - **[no] queue-override**
 - **[no] queue** *queue-id*
 - **adaptation-rule** [*pir adaptation-rule*] [*cir adaptation-rule*]
 - **no adaptation-rule**
 - **avg-frame-overhead** *percentage*
 - **no avg-frame-overhead**
 - **cbs** *size-in-kbytes*
 - **no cbs**
 - **drop-tail** *low*
 - **percent-reduction-from-mbs** *percent*
 - **no percent-reduction-from-mbs**
 - **mbs** *size* {[*bytes* | *kilobytes*] | *default*}
 - **no mbs**
 - **monitor-depth**
 - **[no] monitor-depth**
 - **rate** *pir-rate* [*cir cir-rate*]
 - **no rate**
 - **[no] scheduler-override**
 - **[no] scheduler** *scheduler-name*
 - **parent** [*weight weight*] [*cir-weight cir-weight*]
 - **no parent**
 - **rate** *pir-rate* [*cir cir-rate*]
 - **no rate**
 - **scheduler-policy** *scheduler-policy-name*
 - **no scheduler-policy**
 - **ip-tunnel** *name* [*create*]
 - **no ip-tunnel** *name*
 - **backup-remote-ip** *ip-address*
 - **no backup-remote-ip**
 - **[no] clear-df-bit**
 - **delivery-service** *service-id*
 - **delivery-service** *name service-name*
 - **no delivery-service**
 - **description** *description-string*
 - **no description**

- **dscp** *dscp-name*
- **no dscp**
- **private-tcp-mss-adjust** *octets*
- **private-tcp-mss-adjust** **default**
- **no private-tcp-mss-adjust**
- **public-tcp-mss-adjust** *octets*
- **public-tcp-mss-adjust** **default**
- **remote-ip** *ip-address*
- **no remote-ip**
- **source** *ip-address*
- **no source**
- **lag-link-map-profile** *lag-link-map-profile-id*
- **no lag-link-map-profile**
- **multi-service-site** *customer-site-name*
- **no multi-service-site**
- **[no] shutdown**
- **static-host** **ip** *ip/did-address* [**mac** *ieee-address*] [**create**]
- **static-host** **mac** *ieee-address* [**create**]
- **no static-host** [**ip** *ip-address*] **mac** *ieee-address*
- **no static-host** **all** [**force**]
- **no static-host** **ip** *ip-address*
 - **ancp-string** *ancp-string*
 - **no ancp-string**
 - **app-profile** *app-profile-name*
 - **no app-profile**
 - **inter-dest-id** *intermediate-destination-id*
 - **no inter-dest-id**
 - **[no] shutdown**
 - **sla-profile** *sla-profile-name*
 - **no sla-profile**
 - **sub-profile** *sub-profile-name*
 - **no sub-profile**
 - **subscriber** *sub-ident*
 - **no subscriber**
 - **[no] subscriber-sap-id**
- **transit-policy** {**ip** *ip-aasub-policy-id* | **prefix** *prefix-aasub-policy-id*}
- **no transit-policy**

3.8.1.21 Interface SAP Tunnel Commands

- ```

config
— service
 — vpn service-id [customer customer-id]
 — no vpn service-id
 — [no] interface ip-int-name [create] [tunnel]
 — [no] sap sap-id
 — aarp aarpId type type
 — no aarp
 — accounting-policy acct-policy-id
 — no accounting-policy [acct-policy-id]

```

---

```

— anti-spoof {ip | mac | ip-mac}
— no anti-spoof
— app-profile app-profile-name
— no app-profile
— atm
 — atm
 — ingress traffic-desc-profile-id
 — no ingress
 — encapsulation atm-encap-type
 — ingress
 — ingress traffic-desc-profile-id
 — no ingress
 — oam
 — [no] alarm-cells
 — [no] alarm-cells
— calling-station-id calling-station-id
— no calling-station-id
— [no] app-profile
— cpu-protection policy-id [mac-monitoring] | [eth-cfm-monitoring [aggregate] [car]]
— no cpu-protection
— description description-string
— no description [description-string]
— cpu-protection policy-name
— no cpu-protection
— egress
 — [no] agg-rate
 — [no] rate
 — [no] queue-frame-based-accounting
 — rate {max | rate}
 — no rate
 — [no] hsmda-queue-override
 — secondary-shaper secondary-shaper-name
 — no secondary-shaper
 — wrr-policy hsmda-wrr-policy-name
 — no wrr-policy
 — packet-byte-offset {add add-bytes | subtract sub-bytes}
 — no packet-byte-offset
 — [no] queue queue-id
 — mbs size {[bytes | kilobytes] | default}
 — no mbs
 — rate pir-rate
 — no rate
 — slope-policy hsmda-slope-policy-name
 — no slope-policy
 — wrr-weight weight
 — no wrr-weight
 — source ip-address
 — source ip-address
 — backup-remote-ip ip-address
 — [no] qinq-mark-top-only
 — qos policy-id [port-redirect-group queue-group-name instance instance-id]

```

- **no qos**
- **[no] queue-override**
  - **[no] queue** *queue-id*
    - **adaptation-rule** [*pir adaptation-rule*] [*cir adaptation-rule*]
    - **no adaptation-rule**
    - **adaptation-rule** *percentage*
    - **no adaptation-rule**
    - **cbs** *size-in-kbytes*
    - **no cbs**
    - **drop-tail low**
      - **percent-reduction-from-mbs** *percent*
      - **no percent-reduction-from-mbs**
    - **mbs** *size* {[*bytes* | *kilobytes*] | *default*}
    - **no mbs**
    - **parent** [*weight weight*] [*cir-weight cir-weight*]
    - **no parent**
    - **percent-rate** *pir-percent* [*cir cir-percent*]
    - **no percent-rate**
    - **rate** *pir-rate* [*cir cir-rate*]
    - **no rate**
  - **[no] scheduler-override**
    - **[no] scheduler** *scheduler-name*
      - **parent** [*weight weight*] [*cir-weight cir-weight*]
      - **no parent**
      - **parent** *pir-rate* [*cir cir-rate*]
      - **no parent**
  - **qos** *scheduler-policy-name*
  - **no qos**
- **eth-cfm**
  - **[no] collect-lmm-stats**
  - **collect-lmm-stats** *mep-id domain md-index association ma-index* [*direction* {*up* | *down*}]
  - **no collect-lmm-stats** *mep-id domain md-index association ma-index*
    - **[no] ais-enable**
      - **[no] interface-support-enable**
    - **[no] interface-support-enable**
    - **ccm-ltm-priority** *priority*
    - **no ccm-ltm-priority**
    - **[no] ccm-padding-size** *ccm-padding*
    - **[no] ccm-padding-size**
      - **multiplier** *multiplier-value*
      - **no multiplier**
    - **description** *description-string*
    - **no description**
    - **[no] eth-test-enable**
      - **[no] test-pattern** {*all-zeros* | *all-ones*} [*crc-enable*]
    - **fault-propagation-enable** {*use-if-tlv* | *suspend-ccm*}

- 
- **no fault-propagation-enable**
  - **low-priority-defect** {allDef | macRemErrXcon | remErrXcon | errXcon | xcon | noXcon}
  - **one-way-delay-threshold** *seconds*
  - **squelch-ingress-levels** [*md-level* [*md-level*...]]
  - **no squelch-ingress-levels**
  - **squelch-ingress-levels** [accept | ignore]
  - **frame-relay**
    - [no] **frf-12**
      - **frf-12** *threshold*
      - **no frf-12**
      - [no] **interleave**
    - **scheduling-class** *class-id*
    - **no scheduling-class**
  - [no] **scheduling-class**
  - **host-lockout-policy** *policy-name*
  - **no host-lockout-policy**
  - [no] **host-shutdown**
  - **host-shutdown** *name* [create]
  - **no host-shutdown** *name*
    - **backup-remote-ip** *ip-address*
    - **no backup-remote-ip**
    - [no] **bfd-enable**
    - **delivery-service** *service-id*
    - **no delivery-service**
    - **description** *description-string*
    - **no description**
    - **delivery-service** *dscp-name*
    - **no delivery-service**
    - **remote-ip** *ip-address*
    - **no remote-ip**
    - **source** *ip-address*
    - **no source**
  - **lag-link-map-profile** *lag-link-map-profile-id*
  - **no lag-link-map-profile**
  - **lag-link-map-profile** *customer-site-name*
  - **no lag-link-map-profile**
  - [no] **shutdown**
  - **static-host** **ip** *ip/did-address* [**mac** *ieee-address*] [create]
  - **static-host** **mac** *ieee-address* [create]
  - **no static-host** [**ip** *ip-address*] [**mac** *ieee-address*]
  - **no static-host** **all** [force]
  - **no static-host** **ip** *ip-address*
    - **ancp-string** *ancp-string*
    - **no ancp-string**
    - **app-profile** *app-profile-name*
    - **no app-profile**
    - **app-profile** *intermediate-destination-id*
    - **no app-profile**
    - [no] **shutdown**
    - **route** *sla-profile-name*
    - **no route**
    - **sub-profile** *sub-profile-name*
    - **no sub-profile**

- **subscriber** *sub-ident*
- **no subscriber**
- **[no] subscriber**
- **aarp** *aarpid type type*
- **no aarp**
- **transit-policy ip** *ip-aasub-policy-id*
- **transit-policy prefix** *prefix-aasub-policy-id*
- **no transit-policy**

### 3.8.1.22 Routed VPLS Commands

Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information.

- ```

config
— service
  — vprn service-id [customer customer-id]
  — no vprn service-id
    — interface ip-interface-name [create]
    — no interface ip-interface-name
      — vpls service-name
      — no vpls
        — ingress
          — v4-routed-override-filter ipv4-filter-id
          — no v4-routed-override-filter
          — v6-routed-override-filter ipv6-filter-id
          — no v6-routed-override-filter
        — egress
          — reclassify-using-qos sap-egress-qos-id
          — no reclassify-using-qos
          — v4-routed-override-filter ipv4-filter-id
          — no v4-routed-override-filter
          — v6-routed-override-filter ipv6-filter-id
          — no v6-routed-override-filter

```

3.8.1.23 Oper Group Commands

- ```

config
— service
 — vprn service-id
 — site name [create]
 — monitor-oper-group name
 — no monitor-oper-group name

```

### 3.8.1.24 Network Ingress Commands

```

config
 — service
 — vprn [customer customer-id] [create]
 — no vprn service-id
 — network
 — ingress
 — filter {ip ip-filter-id | ipv6 ipv6-filter-id}
 — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
 — qos network-policy-id fp-redirect-group queue-group-name
 instance instance-id
 — no qos
 — [no] urpf-check

```

### 3.8.1.25 BGP Configuration Commands

```

config
 — service
 — vprn service-id [customer customer-id]
 — no vprn service-id
 — [no] bgp-shared-queue
 — [no] bgp
 — [no] advertise-inactive
 — [no] aggregator-id-zero
 — [no] always-compare-med
 — auth-keychain name
 — authentication-key [authentication-key | hash-key] [hash | hash2]
 — no authentication-key
 — [no] backup-path [ipv4] [label-ipv4] [ipv6]
 — best-path-selection
 — always-compare-med [zero | infinity]
 — always-compare-med strict-as {zero | infinity}
 — no always-compare-med
 — as-path-ignore [ipv4] [label-ipv4] [ipv6]
 — no as-path-ignore
 — [no] deterministic-med
 — ebgp-ibgp-equal [ipv4] [label-ipv4] [ipv6]
 — no ebgp-ibgp-equal
 — [no] ignore-nh-metric
 — [no] ignore-router-id
 — [no] bfd-enable
 — cluster cluster-id
 — no cluster
 — [no] connect-retry seconds
 — [no] damp-peer-oscillations [idle-hold-time initial-wait second-wait
 max-wait] [error-interval minutes]
 — [no] damping
 — description description-string
 — no description

```

- 
- [no] **disable-4byte-asn**
  - [no] **disable-client-reflect**
  - **disable-communities** [standard] [extended] [large]
  - no **disable-communities**
  - [no] **disable-fast-external-failover**
  - **dynamic-neighbor-limit** *peers*
  - no **dynamic-neighbor-limit**
  - [no] **eibgp-loadbalance**
  - **enable-bgp-vpn-backup** [ipv4] [ipv6]
  - no **enable-bgp-vpn-backup**
  - [no] **enable-peer-tracking**
  - [no] **enforce-first-as**
  - **error-handling**
    - [no] **update-fault-tolerance**
  - **export** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]
  - no **export**
  - **family** [ipv4] [label-ipv4] [ipv6] [mcast-ipv4] [flow-ipv4]
  - no **family**
  - **flowspec**
    - [no] **validate-dest-prefix**
    - [no] **validate-redirect-ip**
  - [no] **graceful-restart**
    - [no] **enable-notification**
    - [no] **long-lived**
      - **advertise-stale-to-all-neighbors** [without-no-export | no without-no-export]
      - no **advertise-stale-to-all-neighbors**
      - **advertised-stale-time** *seconds*
      - no **advertised-stale-time**
      - [no] **family** {ipv4 | ipv6 | label-ipv4 | flow-ipv4 | flow-ipv6}
        - **advertised-stale-time** *seconds*
        - no **advertised-stale-time**
        - **helper-override-stale-time** *seconds*
        - no **helper-override-stale-time**
      - **forwarding-bits-set** {all | non-fwd}
      - no **forwarding-bits-set**
      - **helper-override-restart-time** *seconds*
      - no **helper-override-restart-time**
      - **helper-override-stale-time** *seconds*
      - no **helper-override-stale-time**
    - **restart-time** *seconds*
    - no **restart-time**
    - **stale-routes-time** *time*
    - no **stale-routes-time**
  - **hold-time** *seconds* [*min seconds2*]
  - no **hold-time**
  - [no] **ibgp-multipath**
  - **import** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]
  - no **import**
  - [no] **initial-send-delay-zero**
  - **keepalive** *seconds*
  - no **keepalive**
  - **label-preference** *value*

- 
- **no label-preference**
  - **local-as** *as-number* [**private**]
  - **no local-as**
  - **local-preference** *local-preference*
  - **no local-preference**
  - **loop-detect** {**drop-peer** | **discard-route** | **ignore-loop** | **off**}
  - **no loop-detect**
  - **med-out** {*number* | *igp-cost*}
  - **no med-out**
  - **min-route-advertisement** *seconds*
  - **no min-route-advertisement**
  - **multihop** *ttl-value*
  - **no multihop**
  - **multipath** *max-paths* [**ebgp** *ebgp-max-paths*] [**ibgp** *ibgp-max-paths*]  
[**restrict** {**same-neighbor-as** | **exact-as-path**}]
  - **no multipath**
  - **next-hop-resolution**
    - **policy** *policy-name*
  - **peer-tracking-policy** *policy-name*
  - **no peer-tracking-policy**
  - **preference** *preference*
  - **no preference**
  - [**no**] **rapid-withdrawal**
  - **remove-private** [**limited**] [**skip-peer-as**]
  - **no remove-private**
  - **rib-management**
    - **ipv4**
      - **leak-import** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]
      - **no leak-import**
      - **route-table-import** *policy-name*
      - **no route-table-import**
    - **ipv6**
      - **leak-import** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]
      - **no leak-import**
      - **route-table-import** *policy-name*
      - **no route-table-import**
    - **label-ipv4**
      - **leak-import** *plcy-or-long-expr* [*plcy-or-expr*]
      - **no leak-import**
      - **route-table-import** *policy-name*
      - **no route-table-import**
  - **router-id** *ip-address*
  - **no router-id**
  - [**no**] **shutdown**
  - [**no**] **split-horizon**
  - [**no**] **third-party-nexthop**



### 3.8.1.26 BGP Group Configuration Commands

```

config
— service
 — vprn service-id [customer customer-id]
 — no vprn service-id
 — [no] bgp
 — [no] group name [esm-dynamic-peer]
 — [no] advertise-inactive
 — [no] aggregator-id-zero
 — [no] as-override
 — auth-keychain name
 — authentication-key {authentication-key | hash-key} [hash |
 hash2]
 — no authentication-key
 — [no] bfd-enable
 — cluster cluster-id
 — no cluster
 — connect-retry seconds
 — no connect-retry
 — [no] damp-peer-oscillations [idle-hold-time initial-wait
 second-wait max-wait] [error-interval minutes]
 — [no] damping
 — description description-string
 — no description
 — [no] disable-4byte-asn
 — [no] disable-capability-negotiation
 — [no] disable-client-reflect
 — disable-communities [standard] [extended] [large]
 — no disable-communities
 — [no] disable-fast-external-failover
 — dynamic-neighbor
 — [no] prefix ip-prefix/prefix-length
 — dynamic-neighbor-limit peers
 — no dynamic-neighbor-limit
 — [no] enable-peer-tracking
 — [no] enforce-first-as
 — error-handling
 — [no] update-fault-tolerance
 — export policy-name [policy-name]
 — no export
 — family [ipv4] [label-ipv4] [ipv6] [mcast-ipv4] [flow-ipv6] [flow-
 ipv4] [mcast-ipv6]
 — no family
 — [no] graceful-restart
 — [no] enable-notification
 — [no] long-lived
 — advertise-stale-to-all-neighbors [without-no-
 export | no without-no-export]
 — no advertise-stale-to-all-neighbors
 — advertised-stale-time seconds
 — no advertised-stale-time

```

---

```

— [no] family {ipv4 | ipv6 | label-ipv4 | flow-ipv4 |
 flow-ipv6}
 — advertised-stale-time seconds
 — no advertised-stale-time
 — helper-override-stale-time seconds
 — no helper-override-stale-time
 — forwarding-bits-set {all | non-fwd}
 — no forwarding-bits-set
 — helper-override-restart-time seconds
 — no helper-override-restart-time
 — helper-override-stale-time seconds
 — no helper-override-stale-time
— restart-time seconds
— no restart-time
— stale-routes-time time
— no stale-routes-time
— hold-time seconds [min seconds2]
— no hold-time
— import policy-name [policy-name]
— no import
— [no] initial-send-delay-zero
— keepalive seconds
— no keepalive
— label-preference value
— no label-preference
— link-bandwidth
 — accept-from-ebgp family [family]
 — no accept-from-ebgp
 — add-to-received-ebgp family [family]
 — no add-to-received-ebgp
 — aggregate-used-paths family [family]
 — no aggregate-used-paths
 — send-to-ebgp family [family]
 — no send-to-ebgp
— local-address ip-address
— no local-address
— local-as as-number [private] [no-prepend-global-as]
— no local-as
— local-preference local-preference
— no local-preference
— loop-detect {drop-peer | discard-route | ignore-loop | off}
— no loop-detect
— med-out {number | igp-cost}
— no med-out
— min-route-advertisement seconds
— no min-route-advertisement
— multihop tth-value
— no multihop
— [no] next-hop-self
— [no] passive
— peer-as as-number
— no peer-as
— preference preference
— no preference

```

- **prefix-limit** *family limit* [**log-only**] [**threshold percentage**] [**idle-timeout** {*minutes* | **forever**}] [**log-only**] [**post-import**]
- **no prefix-limit** *family*
- [**no**] **remove-private** [**limited**] [**skip-peer-as**]
- [**no**] **shutdown**
- [**no**] **third-party-nexthop**
- **ttl-security** *min-ttl-value*
- **no ttl-security**
- **type** {**internal** | **external**}
- **no type**
- [**no**] **updated-error-handling**

### 3.8.1.27 BGP Group Neighbor Configuration Commands

```

config
 — service
 — vpn service-id [customer customer-id]
 — no vpn service-id
 — [no] bgp
 — [no] group name [esm-dynamic-peer]
 — [no] neighbor ip-address
 — [no] advertise-inactive
 — [no] aggregator-id-zero
 — [no] as-override
 — auth-keychain name
 — authentication-key authentication-key | hash-key [hash | hash2]
 — no authentication-key
 — [no] bfd-enable
 — cluster cluster-id
 — no cluster
 — connect-retry seconds
 — no connect-retry
 — [no] damp-peer-oscillations [idle-hold-time initial-wait second-wait max-wait] [error-interval minutes]
 — [no] damping
 — description description-string
 — no description
 — [no] disable-4byte-asn
 — [no] disable-capability-negotiation
 — [no] disable-client-reflect
 — disable-communities [standard] [extended] [large]
 — no disable-communities
 — [no] disable-fast-external-failover
 — [no] enable-peer-tracking
 — [no] enforce-first-as
 — error-handling
 — [no] update-fault-tolerance
 — export policy-name [policy-name]
 — no export

```

- 
- **family** [ipv4] [label-ipv4] [ipv6] [mcast-ipv4] [flow-ipv6]  
[flow-ipv4] [mcast-ipv6]
  - **no family**
  - [no] **graceful-restart**
    - [no] **enable-notification**
    - [no] **long-lived**
      - **advertise-stale-to-all-neighbors**  
[without-no-export | no without-no-export]
      - **no advertise-stale-to-all-neighbors**
      - **advertised-stale-time** *seconds*
      - **no advertised-stale-time**
      - [no] **family** {ipv4 | ipv6 | label-ipv4 | flow-ipv4 | flow-ipv6}
        - **advertised-stale-time** *seconds*
        - **no advertised-stale-time**
        - **helper-override-stale-time**  
*seconds*
        - **no helper-override-stale-time**
      - **forwarding-bits-set** {all | non-fwd}
      - **no forwarding-bits-set**
      - **helper-override-restart-time** *seconds*
      - **no helper-override-restart-time**
      - **helper-override-stale-time** *seconds*
      - **no helper-override-stale-time**
    - **restart-time** *seconds*
    - **no restart-time**
    - **stale-routes-time** *time*
    - **no stale-routes-time**
  - **hold-time** *seconds* [min *seconds2*]
  - **no hold-time**
  - **import** *policy-name* [*policy-name*]
  - **no import**
  - [no] **initial-send-delay-zero**
  - **keepalive** *seconds*
  - **no keepalive**
  - **label-preference** *value*
  - **no label-preference**
  - **link-bandwidth**
    - **accept-from-ebgp** *family* [*family*]
    - **no accept-from-ebgp**
    - **add-to-received-ebgp** *family* [*family*]
    - **no add-to-received-ebgp**
    - **aggregate-used-paths** *family* [*family*]
    - **no aggregate-used-paths**
    - **send-to-ebgp** *family* [*family*]
    - **no send-to-ebgp**
  - **local-address** *ip-address*
  - **no local-address**
  - **local-as** *as-number* [private] [no-prepend-global-as]
  - **no local-as**
  - **local-preference** *local-preference*
  - **no local-preference**

- **loop-detect** {drop-peer | discard-route | ignore-loop | off}
- **no loop-detect**
- **med-out** {number | igp-cost}
- **no med-out**
- **min-route-advertisement** seconds
- **no min-route-advertisement**
- **multihop** ttl-value
- **no multihop**
- [no] **next-hop-self**
- [no] **passive**
- **peer-as** as-number
- **no peer-as**
- **preference** preference
- **no preference**
- **prefix-limit** family limit [log-only] [threshold percentage] [idle-timeout {minutes | forever}] [log-only] [post-import]
- **no prefix-limit** family
- [no] **remove-private** [limited] [skip-peer-as]
- [no] **shutdown**
- [no] **third-party-nexthop**
- **ttl-security** min-ttl-value
- **no ttl-security**
- **type** {internal | external}
- **no type**
- [no] **updated-error-handling**

### 3.8.1.28 IS-IS Configuration Commands

- ```

config
  — service
    — vprn service-id [customer customer-id]
    — no vprn service-id
      — [no] isis [isis-instance]
        — [no] advertise-passive-only
        — advertise-router-capability {area | as}
        — no advertise-router-capability
        — all-l1isis ieee-address
        — no all-l1isis
        — all-l2isis ieee-address
        — no all-l2isis
        — [no] area-id area-address
        — auth-keychain name
        — [no] authentication-check
        — authentication-key [authentication-key | hash-key] [hash | hash2]
        — no authentication-key
        — authentication-type {password | message-digest}
        — no authentication-type
        — [no] csnp-authentication
        — default-route-tag tag
  
```

```

— no default-route-tag
— export policy-name [policy-name]
— no export
— export-limit number [log percentage]
— no export-limit
— [no] graceful-restart
    — [no] helper-disable
— [no] hello-authentication
— hello-padding {none | adaptive | loose | strict}
— no hello-padding
— [no] ignore-attached-bit
— [no] ignore-lsp-errors
— [no] iid-tlv-enable
— import policy-name [policy-name]
— no import
— [no] interface ip-int-name
    — bfd-enable {ipv4 | ipv6} [include-bfd-tlv]
    — no bfd-enable
    — csnp-interval seconds
    — no csnp-interval
    — [no] default-instance
    — hello-auth-keychain name
    — no hello-auth-keychain
    — [no] hello-authentication
    — hello-authentication-key authentication-key | hash-key [hash |
        hash2]
    — no hello-authentication-key
    — hello-authentication-type {password | message-digest}
    — no hello-authentication-type
    — hello-padding {none | adaptive | loose | strict}
    — no hello-padding
    — interface-type {broadcast | point-to-point}
    — no interface-type
    — [no] ipv4-multicast-disable
    — [no] ipv6-unicast-disable
    — level {1 | 2}
        — hello-auth-keychain name
        — no hello-auth-keychain
        — hello-authentication-key authentication-key | hash-key
            [hash | hash2]
        — no hello-authentication-key
        — hello-authentication-type {password | message-
            digest}
        — no hello-authentication-type
        — hello-interval seconds
        — no hello-interval
        — hello-multiplier multiplier
        — no hello-multiplier
        — hello-padding {none | adaptive | loose | strict}
        — no hello-padding
        — ipv4-multicast-metric IPv4 multicast metric
        — no ipv4-multicast-metric
        — ipv6-unicast-metric ipv6-metric
        — no ipv6-unicast-metric

```

-
- **metric** *ipv4-metric*
 - **no metric**
 - **[no] passive**
 - **priority** *number*
 - **no priority**
 - **sd-offset** *sd-offset*
 - **no sd-offset**
 - **sf-offset** *sf-offset*
 - **no sf-offset**
 - **level-capability** {*level-1* | *level-2* | *level-1/2*}
 - **no level-capability**
 - **lfa-policy-map** *route-nh-template* *template-name*
 - **no lfa-policy-map**
 - **load-balancing-weight** *weight*
 - **no load-balancing-weight**
 - **[no] loopfree-alternate-exclude**
 - **lsp-pacing-interval** *milliseconds*
 - **no lsp-pacing-interval**
 - **mesh-group** [*value* | **blocked**]
 - **no mesh-group**
 - **[no] passive**
 - **retransmit-interval** *seconds*
 - **no retransmit-interval**
 - **[no] shutdown**
 - **tag** *tag*
 - **no tag**
 - **ipv4-multicast-routing** {*native* | *mt*}
 - **no ipv4-multicast-routing**
 - **[no] ipv4-routing**
 - **ipv6-routing** {*native* | *mt*}
 - **no ipv6-routing**
 - **level** *level*
 - **[no] advertise-router-capability**
 - **auth-keychain** *name*
 - **authentication-key** *authentication-key* | *hash-key* [**hash** | **hash2**]
 - **no authentication-key**
 - **authentication-type** {*password* | *message-digest*}
 - **no authentication-type**
 - **[no] csnp-authentication**
 - **default-ipv4-multicast-metric** *ipv4 multicast metric*
 - **no default-ipv4-multicast-metric**
 - **default-ipv6-multicast-metric**
 - **no default-ipv6-multicast-metric**
 - **default-ipv6-unicast-metric** *ipv6 metric*
 - **no default-ipv6-unicast-metric**
 - **default-metric** *ipv4 metric*
 - **no default-metric**
 - **external-preference** *external-preference*
 - **no external-preference**
 - **[no] hello-authentication**
 - **hello-padding** {*none* | *adaptive* | *loose* | *strict*}
 - **no hello-padding**
 - **[no] loopfree-alternate-exclude**

-
- **lsp-mtu-size** *size*
 - **no lsp-mtu-size**
 - **preference** *preference*
 - **no preference**
 - **[no] wide-metrics-only**
 - **level-capability** {*level-1* | *level-2* | *level-1/2*}
 - **[no] link-group** *link-group name*
 - **description** [*256 chars max*]
 - **no description**
 - **level** {*1* | *2*}
 - **ipv4-multicast-metric-offset** *offset-value*
 - **no ipv4-multicast-metric-offset**
 - **ipv4-unicast-metric-offset** *offset-value*
 - **no ipv4-unicast-metric-offset**
 - **ipv6-unicast-metric-offset** *offset-value*
 - **no ipv6-unicast-metric-offset**
 - **[no] member** *interface-name*
 - **oper-members** *oper-members*
 - **no oper-members**
 - **revert-members** *revert-members*
 - **no revert-members**
 - **[no] loopfree-alternates**
 - **exclude**
 - **prefix-policy** *prefix-policy* [*prefix-policy*]
 - **no prefix-policy**
 - **lsp-lifetime** *seconds*
 - **no lsp-lifetime**
 - **lsp-mtu-size** *size*
 - **no lsp-mtu-size**
 - **lsp-refresh-interval** [*seconds*] [*half-lifetime* {*enable* | *disable*}]
 - **no lsp-refresh-interval**
 - **[no] multi-topology**
 - **[no] ipv4-multicast**
 - **[no] ipv6-unicast**
 - **[no] multicast-import** [*ipv4*]
 - **overload** [*timeout seconds*] [*max-metric*]
 - **no overload**
 - **[no] overload-export-external**
 - **[no] overload-export-interlevel**
 - **overload-on-boot** [*timeout seconds*] [*max-metric*]
 - **no overload-on-boot**
 - **[no] poi-tlv-enable**
 - **[no] prefix-attributes-tlv**
 - **prefix-limit** *limit* [*log-only*] [*threshold percent*] [*overload-timeout* {*seconds* | *forever*}]
 - **[no] psnp-authentication**
 - **no prefix-limit**
 - **reference-bandwidth** *bandwidth-in-kbps*
 - **reference-bandwidth** [*tbps Tera-bps*] [*gbps Giga-bps*] [*mbps Mega-bps*] [*kbps Kilo-bps*]
 - **no reference-bandwidth**
 - **rib-priority** {*high*} *prefix-list-name* | **tag** *tag-value*
 - **no rib-priority**
 - **router-id** *router-id*

- **no router-id**
- **[no] shutdown**
- **[no] standard-multi-instance**
- **[no] strict-adjacency-check**
- **summary-address** {*ip-prefix/mask* | *ip-prefix [netmask]*} [*level*] [**tag** *tag*]
- **no summary-address** {*ip-prefix/mask* | *ip-prefix [netmask]*}
- **[no] suppress-attached-bit**
- **system-id** *isis-system-id*
- **no system-id**
- **[no] timers**
 - **lsp-wait** *lsp-wait* [**lsp-initial-wait** *initial-wait*] [**lsp-second-wait** *second wait*]
 - **no lsp-wait**
 - **spf-wait** *spf-wait* [**spf-initial-wait** *initial-wait*] [**spf-second-wait** *second wait*]
 - **no spf-wait**
- **[no] unicast-import-disable**

3.8.1.29 OSPF Configuration Commands

- ```

config
 — service
 — vpn service-id [customer customer-id]
 — no vpn service-id
 — ospf [router-id]
 — no ospf
 — advertise-router-capability {link | area | as}
 — no advertise-router-capability
 — [no] area area-id
 — [no] advertise-router-capability
 — area-range ip-prefix/mask [advertise | not-advertise]
 — no area-range ip-prefix/mask
 — [no] blackhole-aggregate
 — export policy-name [policy-name]
 — no export
 — import policy-name [policy-name]
 — no import
 — interface ip-int-name [secondary]
 — no interface ip-int-name
 — [no] advertise-router-capability
 — [no] advertise-subnet
 — auth-keychain name
 — 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 | non-
 broadcast}
— no interface-type
— lfa-policy-map route-nh-template template-name
— no lfa-policy-map
— [no] loopfree-alternate-exclude
— lsa-filter-out [all | except-own-rtrlsa | except-own-
 rtrlsa-and-defaults]
— no lsa-filter-out
— 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] neighbor ip-address
— [no] passive
— poll-interval seconds
— no poll-interval
— priority number
— no priority
— retransmit-interval seconds
— no retransmit-interval
— rib-priority prefix-list-name
— no rib-priority
— [no] shutdown
— transit-delay seconds
— no transit-delay
— [no] loopfree-alternate-exclude
— [no] nssa
 — area-range ip-prefix/mask [advertise | not-advertise]
 — no area-range ip-prefix/mask
 — originate-default-route [type-7] [no-adjacency-check]
 — no originate-default-route
 — [no] redistribute-external
 — [no] summaries
— [no] sham-link ip-int-name ip-address
 — auth-keychain name
 — 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*
    - **auth-keychain** *name*
    - **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*]
  - **no export**
  - **export-limit** *number* [*log percentage*]
  - **no export-limit**
  - **external-db-overflow** *limit seconds*
  - **no external-db-overflow**
  - **external-preference** *preference*
  - **no external-preference**
  - **[no] graceful-restart**
    - **[no] helper-disable**
    - **[no] strict-lsa-checking**
  - **[no] ignore-dn-bit**
  - **import** *policy-name* [*policy-name*]
  - **no import**
  - **[no] ignore-dn-bit**
  - **[no] loopfree-alternates**
    - **exclude**
      - **prefix-policy** *prefix-policy* [*prefix-policy*]
      - **no prefix-policy**
  - **[no] loopfree-alternates**
  - **[no] multicast-import**
  - **overload** [*timeout seconds*]
  - **no overload**
  - **[no] overload-include-ext-2**

---

```

— [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
— rib-priority prefix-list-name
— no rib-priority
— router-id ip-address
— no router-id
— rtr-adv-lsa-limit [1..4294967295] [log-only] [threshold percent]
— rtr-adv-lsa-limit [1..4294967295] [log-only] [threshold percent]
 overload-timeout forever
— rtr-adv-lsa-limit [1..4294967295] [log-only] [threshold percent]
 overload-timeout seconds
— no rtr-adv-lsa-limit
— [no] shutdown
— [no] super-backbone
— [no] suppress-dn-bit
— timers
 — incremental-spf-wait inc-spf-wait
 — no incremental-spf-wait
 — lsa-accumulate lsa-accum-time
 — no lsa-accumulate
 — lsa-arrival lsa-arrival-time
 — no lsa-arrival
 — lsa-generate max-lsa-wait [lsa-initial-wait lsa-initial-wait [lsa-
 second-wait lsa-second-wait]]
 — no lsa-generate
 — redistribute-delay redistrib-wait
 — no redistribute-delay
 — spf-wait max-spf-wait [spf-initial-wait [spf-second-wait]]
 — no spf-wait
— [no] unicast-import-disable
— vpn-domain id {0005 | 0105 | 0205 | 8005}
— no vpn-domain
— vpn-tag vpn-tag
— no vpn-tag
— ospf3 [instance-id] [router-id]
— [no] ospf3 instance-id
 — advertise-router-capability {link | area | as}
 — no advertise-router-capability
 — [no] area area-id
 — [no] advertise-router-capability
 — area-range ip-prefix/mask [advertise | not-advertise]
 — no area-range ip-prefix/mask
 — [no] blackhole-aggregate
 — export policy-name [policy-name]
 — no export
 — import policy-name [policy-name]
 — no import
 — interface ip-int-name [secondary]
 — no interface ip-int-name

```

- 
- [no] **advertise-router-capability**
  - **authentication** *bidirectional sa-name*
  - **authentication** *inbound sa-name outbound sa-name*
  - **no authentication**
  - **bfd-enable** [*remain-down-on-failure*]
  - **no bfd-enable**
  - **dead-interval** *seconds*
  - **no dead-interval**
  - **hello-interval** *seconds*
  - **no hello-interval**
  - **interface-type** {*broadcast | point-to-point | non-broadcast*}
  - **no interface-type**
  - **lfa-policy-map** *route-nh-template template-name*
  - **no lfa-policy-map**
  - [no] **loopfree-alternate-exclude**
  - **lsa-filter-out** [*all | except-own-rtlsa | except-own-rtlsa-and-defaults*]
  - **no lsa-filter-out**
  - **metric** *metric*
  - **no metric**
  - **mtu** *bytes*
  - **no mtu**
  - [no] **neighbor** *ip-address*
  - [no] **passive**
  - **poll-interval** *seconds*
  - **no poll-interval**
  - **priority** *number*
  - **no priority**
  - **retransmit-interval** *seconds*
  - **no retransmit-interval**
  - **rib-priority** *prefix-list-name*
  - **no rib-priority**
  - [no] **shutdown**
  - **transit-delay** *seconds*
  - **no transit-delay**
  - **key-rollover-interval** *key-rollover-interval*
  - [no] **loopfree-alternate-exclude**
  - [no] **nssa**
    - **area-range** {*ip-prefix/mask | ipv6-prefix/prefix-length*} [*advertise | not-advertise*]
    - **no area-range** {*ip-prefix/mask | ipv6-prefix/prefix-length*}
    - **originate-default-route** [*type-7*] [*no-adjacency-check*]
    - **no originate-default-route**
    - [no] **redistribute-external**
    - [no] **summaries**
  - [no] **stub**
    - **default-metric** *metric*
    - **no default-metric**
    - [no] **summaries**
  - [no] **virtual-link** *router-id transit-area area-id*
    - **authentication** *bidirectional sa-name*
    - **authentication** *inbound sa-name outbound sa-name*
    - **no authentication**

---

```

— dead-interval seconds
— no dead-interval
— hello-interval seconds
— no hello-interval
— retransmit-interval seconds
— no retransmit-interval
— [no] shutdown
— transit-delay seconds
— no transit-delay
— export policy-name [policy-name]
— no export
— export-limit number [log percentage]
— no export-limit
— external-db-overflow limit seconds
— no external-db-overflow
— external-preference preference
— no external-preference
— [no] graceful-restart
 — [no] helper-disable
 — [no] strict-lsa-checking
— [no] ignore-dn-bit
— import policy-name [policy-name]
— no import
— [no] loopfree-alternates
 — exclude
 — prefix-policy prefix-policy prefix-policy [prefix-policy]
 — no prefix-policy
— [no] multicast-import
— overload [timeout seconds]
— no overload
— [no] overload-include-ext-2
— [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
— rib-priority prefix-list-name
— no rib-priority
— router-id ip-address
— no router-id
— rtr-adv-lsa-limit [1..4294967295] [log-only] [threshold percent]
— rtr-adv-lsa-limit [1..4294967295] [log-only] [threshold percent]
 overload-timeout forever
— rtr-adv-lsa-limit [1..4294967295] [log-only] [threshold percent]
 overload-timeout seconds
— no rtr-adv-lsa-limit
— [no] shutdown
— [no] suppress-dn-bit
— timers
 — incremental-spf-wait inc-spf-wait
 — no incremental-spf-wait
 — lsa-accumulate lsa-accum-time

```

- **no lsa-accumulate**
- **lsa-arrival** *lsa-arrival-time*
- **no lsa-arrival**
- **lsa-generate** *max-lsa-wait* [**lsa-initial-wait** *lsa-initial-wait* [**lsa-second-wait** *lsa-second-wait*]]
- **no lsa-generate**
- **redistribute-delay** *redist-wait*
- **no redistribute-delay**
- **spf-wait** *max-spf-wait* [**spf-initial-wait** *spf-initial-wait*] [**spf-second-wait** *spf-second-wait*]
- **no spf-wait**
- **[no] unicast-import-disable**

### 3.8.1.30 PIM Configuration Commands

- ```

config
  — service
    — vrn service-id [customer customer-id]
    — no vrn service-id
    — [no] pim
      — apply-to {all | none}
      — [no] grt-extranet
        — group-prefix ip-address/mask [ip-address/mask] [starg]
        — group-prefix any
        — no group-prefix ip-address/mask
        — no group-prefix any
      — import {join-policy | register-policy} policy-name [policy-name]
      — no import {join-policy | register-policy}
      — [no] interface ip-int-name
        — assert-period assert-period
        — no assert-period
        — [no] bfd-enable {ipv4 | ipv6}
        — [no] bsm-check-rtr-alert
        — hello-interval hello-interval
        — no hello-interval
        — hello-multiplier deci-units
        — no hello-multiplier
        — [no] improved-assert
        — [no] instant-prune-echo
        — [no] ipv4-multicast-disable
        — [no] ipv6-multicast-disable
        — max-groups value
        — no max-groups
        — mcac
          — if-policy mcac-if-policy-name
          — no if-policy
          — mc-constraints
            — level level-id bw bandwidth
            — no level level-id
            — number-down number-lag-port-down level level-id

```

```

— no number-down number-lag-port-down
— [no] shutdown
— [no] use-lag-port-weight
— policy policy-name
— no policy
— unconstrained-bw bandwidth mandatory-bw
   mandatory-bw
— no unconstrained-bw
— monitor-oper-group group-name family {ipv4 | ipv6} {add | set
  | subtract} value
— no monitor-oper-group [family {ipv4 | ipv6}]
— multicast-senders {auto | always | never}
— no multicast-senders
— [no] p2mp-ldp-tree-join [ipv4] [ipv6]
— priority dr-priority
— no priority
— [no] shutdown
— sticky-dr [priority dr-priority]
— no sticky-dr
— three-way-hello [compatibility-mode]
— no three-way-hello
— [no] tracking-support
— [no] ipv4-multicast-disable
— [no] ipv6-multicast-disable
— [no] mc-ecmp-balance
— mc-ecmp-balance-hold minutes
— no mc-ecmp-balance-hold
— mc-ecmp-hashing-enabled [rebalance]
— no mc-ecmp-hashing-enabled
— [no] non-dr-attract-traffic
— rp
  — [no] anycast rp-ip-address
    — [no] rp-set-peer ip-address
  — [no] auto-rp-discovery
  — bootstrap-export policy-name [policy-name]
  — no bootstrap-export
  — bootstrap-import policy-name [policy-name]
  — no bootstrap-import
  — bsr-candidate
    — address ip-address
    — no address
    — hash-mask-len hash-mask-length
    — no hash-mask-len
    — priority bootstrap-priority
    — no priority
    — [no] shutdown
  — ipv6
    — [no] anycast ipv6-address
      — [no] rp-set-peer ipv6-address
    — bsr-candidate
      — address ipv6-address
      — [no] address
      — hash-mask-len hash-mask-length
      — [no] hash-mask-len

```


- **priority** *bootstrap-priority*
- **no priority**
- **[no] shutdown**
- **[no] embedded-rp**
 - **[no] group-range** *grp-ipv6-address/prefix-length*
 - **[no] shutdown**
- **rp-candidate**
 - **address** *ipv6-address*
 - **no address**
 - **[no] group-range** *grp-ipv6-address/prefix-length*
 - **holdtime** *holdtime*
 - **no holdtime**
 - **priority** *priority*
 - **no priority**
 - **[no] shutdown**
- **static**
 - **[no] address** *ipv6-address*
 - **[no] group-prefix** *grp-ipv6-address/prefix-length*
 - **[no] override**
- **rp-candidate**
 - **address** *ip-address*
 - **no address**
 - **[no] group-range** *{grp-ip-address/mask | grp-ip-address [netmask]}*
 - **holdtime** *holdtime*
 - **no holdtime**
 - **priority** *priority*
 - **no priority**
 - **[no] shutdown**
- **static**
 - **[no] address** *ip-address*
 - **[no] group-prefix** *{grp-ip-address/mask | grp-ip-address netmask}*
 - **[no] override**
- **rpf-table** *{rtable-m | rtable-u | both}*
- **no rpf-table**
- **rpf6-table** *{rtable6-m | rtable6-u | both}*
- **no rpf6-table**
- **[no] shutdown**
- **spt-switchover-threshold** *{grp-ip-address/mask | grp-ip-address netmask} spt-threshold*
- **no spt-switchover-threshold** *{grp-ip-address/mask | grp-ip-address netmask}*
- **ssm-assert-compatible-mode** *[enable | disable]*
- **ssm-default-range-disable** *ipv4*
- **[no] ssm-groups**
 - **[no] group-range** *{grp-ip-address/mask | grp-ip-address netmask}*

3.8.1.31 MSDP Configuration Commands

```

config
  — service
    — vprn service-id [customer customer-id]
    — no vprn service-id
      — [no] msdp
        — active-source-limit number
        — no active-source-limit
        — [no] data-encapsulation
        — export policy-name [policy-name]
        — no export
        — [no] group group-name
          — active-source-limit number
          — no active-source-limit
          — export policy-name [policy-name]
          — no export
          — import policy-name [policy-name]
          — no import
          — local-address address
          — no local-address
          — mode {mesh-group | standard}
          — [no] peer peer-address
            — active-source-limit number
            — no active-source-limit
            — authentication-key [authentication-key | hash-key]
              [hash | hash2]
            — no authentication-key
            — [no] default-peer
            — export policy-name [policy-name]
            — no export
            — import policy-name [policy-name]
            — no import
            — local-address address
            — no local-address
            — receive-msdp-msg-rate number interval seconds
              [threshold number]
            — no receive-msdp-msg-rate
            — [no] shutdown
          — receive-msdp-msg-rate number interval seconds [threshold
            number]
          — no receive-msdp-msg-rate
          — [no] shutdown
        — import policy-name [policy-name]
        — no import
        — local-address address
        — no local-address
        — [no] peer peer-address
          — active-source-limit number
          — no active-source-limit
          — authentication-key [authentication-key | hash-key] [hash |
            hash2]
          — no authentication-key

```

- [no] **default-peer**
- **export** *policy-name* [*policy-name*]
- **no export**
- **import** *policy-name* [*policy-name*]
- **no import**
- **local-address** *address*
- **no local-address**
- **receive-msdp-msg-rate** *number* *interval* *seconds* [*threshold* *number*]
- **no receive-msdp-msg-rate**
- [no] **shutdown**
- **receive-msdp-msg-rate** *number* *interval* *seconds* [*threshold* *number*]
- **no receive-msdp-msg-rate**
- **rpf-table** {*rtable-m* | *rtable-u* | *both*}
- **no rpf-table**
- **sa-timeout** *seconds*
- **no sa-timeout**
- [no] **shutdown**
- [no] **source** *prefix/mask*
 - **active-source-limit** *number*
 - **no active-source-limit**

3.8.1.32 MLD Configuration Commands

- ```

config
— service
 — vprn service-id [customer customer-id]
 — no vprn service-id
 — [no] mld
 — [no] group-interface ip-int-name
 — mcac
 — if-policy mcac-if-policy-name
 — no if-policy
 — mc-constraints
 — [no] shutdown
 — policy policy-name
 — no policy
 — unconstrained-bw bandwidth mandatory-bw
 mandatory-bw
 — no unconstrained-bw
 — [no] interface ip-int-name
 — [no] disable-router-alert-check
 — import policy-name
 — no import
 — max-groups value
 — no max-groups
 — mcac
 — if-policy mcac-if-policy-name
 — no if-policy
 — mc-constraints
 — level level-id bw bandwidth

```

- **no level** *level-id*
- **number-down** *number-lag-port-down level level-id*
- **no number-down**
- **[no] shutdown**
- **[no] use-lag-port-weight**
- **policy** *policy-name*
- **no policy**
- **unconstrained-bw** *bandwidth mandatory-bw mandatory-bw*
- **no unconstrained-bw**
- **query-interval** *seconds*
- **no query-interval**
- **query-last-member-interval** *seconds*
- **no query-last-member-interval**
- **query-response-interval** *seconds*
- **no query-response-interval**
- **[no] shutdown**
- **static**
  - **[no] group** *grp-ipv6-address*
  - **[no] group start** *grp-ipv6-address end grp-ipv6-address [step ipv6-address]*
    - **[no] source** *src-ipv6-address*
    - **[no] starg**
- **version** *version*
- **no version**
- **query-interval** *seconds*
- **no query-interval**
- **query-last-member-interval** *seconds*
- **no query-last-member-interval**
- **query-response-interval** *seconds*
- **no query-response-interval**
- **robust-count** *robust-count*
- **no robust-count**
- **[no] shutdown**
- **ssm-translate**
  - **[no] grp-range** *start end*
    - **[no] source** *src-ipv6-address*

### 3.8.1.33 RIP Configuration Commands

- ```

config
  — service
    — vprn service-id [customer customer-id]
    — no vprn service-id
      — [no] rip
      — [no] ripng
        — authentication-key [authentication-key | hash-key] [hash | hash2]
        — no authentication-key
        — authentication-type {none | password | message-digest}
        — no authentication-type

```

-
- [no] **bfd-enable**
 - **check-zero** {enable | disable}
 - no **check-zero**
 - **description** *description-string*
 - no **description**
 - **export** *policy-name* [*policy-name*]
 - no **export**
 - **export-limit** *number* [**log** *percentage*]
 - no **export-limit**
 - [no] **group** *name*
 - **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
 - no **authentication-key**
 - **authentication-type** {none | password | message-digest}
 - no **authentication-type**
 - **check-zero** {enable | disable}
 - no **check-zero**
 - **description** *description-string*
 - no **description**
 - **export** *policy-name* [*policy-name*]
 - no **export**
 - **export-limit** *number* [**log** *percentage*]
 - no **export-limit**
 - **import** *policy-name* [*policy-name*]
 - no **import**
 - **message-size** *max-num-of-routes*
 - no **message-size**
 - **metric-in** *metric*
 - no **metric-in**
 - **metric-out** *metric*
 - no **metric-out**
 - [no] **neighbor** *ip-int-name*
 - **authentication-key** *authentication-key* | *hash-key* [**hash** | **hash2**]
 - no **authentication-key**
 - **authentication-type** {none | password | message-digest}
 - no **authentication-type**
 - **check-zero** {enable | disable}
 - no **check-zero**
 - **description** *description-string*
 - no **description**
 - **export** *policy-name* [*policy-name*]
 - no **export**
 - **export-limit** *number* [**log** *percentage*]
 - no **export-limit**
 - **import** *policy-name* [*policy-name*]
 - no **import**
 - **message-size** *max-num-of-routes*
 - no **message-size**
 - **metric-in** *metric*
 - no **metric-in**
 - **metric-out** *metric*
 - no **metric-out**

- **preference** *preference*
- **no preference**
- **receive** *receive-type*
- **no receive**
- **send** *send-type*
- **no send**
- **[no] shutdown**
- **split-horizon** {*enable* | *disable*}
- **no split-horizon**
- **timers** *update timeout flush*
- **no timers**
- **[no] unicast-address** *ipv6-address*
- **preference** *preference*
- **no preference**
- **receive** *receive-type*
- **no receive**
- **send** *send-type*
- **no send**
- **[no] shutdown**
- **split-horizon** {*enable* | *disable*}
- **no split-horizon**
- **timers** *update timeout flush*
- **no timers**
- **import** *policy-name* [*policy-name*]
- **no import**
- **message-size** *max-num-of-routes*
- **no message-size**
- **metric-in** *metric*
- **no metric-in**
- **metric-out** *metric*
- **no metric-out**
- **preference** *preference*
- **no preference**
- **[no] propagate-metric**
- **receive** *receive-type*
- **no receive**
- **send** *send-type*
- **no send**
- **[no] shutdown**
- **split-horizon** {*enable* | *disable*}
- **no split-horizon**
- **timers** *update timeout flush*
- **no timers**

3.8.1.34 RADIUS Commands

- ```

config
— service
 — vpn service-id [customer customer-id]
 — no vpn service-id
 — radius-proxy

```

- 
- **server** *server-name* [**create**] [**purpose** {[**accounting**] [**authentication**]}]
  - **no server** *server-name*
    - **cache**
      - **key** *packet-type* {**accept** | **request**} *attribute-type* *attribute-type* [*vendor-id* *vendor-id*]
      - **no key**
      - [**no**] **shutdown**
      - **timeout** [*hrs* *hours*] [*min* *minutes*] [*sec* *seconds*]
      - **no timeout**
      - **track-accounting** [**start**] [**stop**] [**interim-update**]
      - **no track-accounting**
    - **default-accounting-server-policy** *policy-name*
    - **no default-accounting-server-policy**
    - **default-authentication-server-policy** *policy-name*
    - **no default-authentication-server-policy**
    - **description** *description-string*
    - **no description**
    - [**no**] **interface** *ip-int-name*
    - **load-balance-key** *vendor* *vendor-id* [*vendor-id*] *attribute-type* *attribute-type* [*attribute-type*]
    - **load-balance-key** *source-ip-udp*
    - **no load-balance-key**
    - **python-policy** *name*
    - **no python-policy**
    - **secret** *secret* [**hash** | **hash2**]
    - **no secret**
    - [**no**] **send-accounting-response**
    - [**no**] **shutdown**
    - **username** *user-name* *prefix-string* [128 chars max] [*accounting-server-policy* *policy-name*] [*authentication-server-policy* *policy-name*]
    - **no username** *user-name*
  - **radius-server**
    - **server** *server-name* [**address** *ip-address*] [**secret** *key*] [**hash** | **hash2**] [**port** *port*] [**create**]
    - **no server** *server-name*
      - [**no**] **accept-coa**
      - **coa-script-policy** *script-policy-name*
      - **no coa-script-policy**
      - **description** *description-string*
      - **no description**
      - **pending-requests-limit** *limit*
      - **no pending-requests-limit**
      - **python-policy** *name*
      - **no python-policy**

### 3.8.1.35 Web Portal Protocol Configuration Commands

- ```
config
— service
```

- **vprn** *service-id* [**customer** *customer-id*]
- **no vprn** *service-id*
 - **[no] wpp**
 - **portals**
 - **portal** *name* **address** *ip-address* [**create**]
 - **[no] portal** *name*
 - **[no] shutdown**
 - **[no] shutdown**

3.8.1.36 AARP Interface Commands

- ```
config
— service
 — vprn service-id [customer customer-id]
```
- **no vprn** *service-id*
    - **aarp-interface** *arp-int-name* [**create**]
    - **no aarp-interface** *arp-int-name*
      - **description** *long-description-string*
      - **no description**
      - **ip-mtu** *octets*
      - **no ip-mtu**
      - **[no] shutdown**
      - **spoke-sdp** *sdp-id:vc-id* [**create**]
      - **no spoke-sdp** *sdp-id:vc-id*
        - **aarp** *aarp-id* **type** {**subscriber-side-shunt** | **network-side-shunt**}
        - **no aarp**
        - **description** *description-string*
        - **no description**
        - **egress**
          - **filter** **ip** *ip-filter-id*
          - **no filter**
          - **vc-label** *vc-label*
          - **no vc-label** [*vc-label*]
        - **ingress**
          - **filter** **ip** *ip-filter-id*
          - **no filter**
          - **vc-label** *vc-label*
          - **no vc-label** [*vc-label*]
      - **[no] shutdown**



## 3.8.2 Command Descriptions

### 3.8.2.1 Generic Commands

#### shutdown

**Syntax** [no] shutdown

**Context**

```

config>service>vprn
config>service>vprn>aaa>remote-servers>radius
config>service>vprn>aarp-interface
config>service>vprn>aarp-interface>spoke-sdp
config>service>vprn>dhcp6>server>failover
config>service>vprn>igmp-trk
config>service>vprn>red-if
config>service>vprn>router-advert>if
config>service>vprn>gsmp
config>service>vprn>gsmp>group
config>service>vprn>gsmp>group>neighbor
config>service>vprn>igmp
config>service>vprn>igmp>grp-if>mcac>mc-constraints
config>service>vprn>igmp>if
config>service>vprn>igmp>if>mcac
config>service>vprn>igmp>if>mcac>mc-constraints
config>service>vprn>if
config>service>vprn>if>dhcp
config>service>vprn>if>dhcp>proxy
config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp
config>service>vprn>if>sap
config>service>vprn>if>sap>static-host
config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
config>service>vprn>isis
config>service>vprn>isis>if
config>service>vprn>mld>grp-if>mcac>mc-constraints
config>service>vprn>mld>interface>mcac>mc-constraints
config>service>vprn>msdp
config>service>vprn>msdp>group
config>service>vprn>msdp>group>peer
config>service>vprn>msdp>peer
config>service>vprn>mvpn>provider-tunnel>inclusive>pim
config>service>vprn>ospf
config>service>vprn>ospf>area>if

```

```

config>service>vprn>ospf3
config>service>vprn>ospf3>area>if
config>service>vprn>ospf3>area>virtual-link
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link
config>service>vprn>red-if>spoke-sdp
config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor
config>service>vprn>pim
config>service>vprn>pim>if
config>service>vprn>pim>if>mcac>mc-constraints
config>service>vprn>pim>rp>bsr-candidate
config>service>vprn>pim>rp>ipv6>bsr-candidate
config>service>vprn>pim>rp>ipv6>embedded-rp
config>service>vprn>pim>rp>ipv6>rp-candidate
config>service>vprn>sub-if>grp-if
config>service>vprn>sub-if>grp-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp>proxy-server
config>service>vprn>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>arp-host
config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
config>service>vprn>sub-if>grp-if>wpp
config>service>vprn>dhcp>server>failover
config>service>vprn>nw-if>dhcp
config>service>vprn>nw-if>eth-cfm>mep
config>service>vprn>radius-proxy>server>cache
config>service>vprn>radius-proxy>server
config>service>vprn>radius-server
config>service>vprn>if>sap>ipsec-tunnel
config>service>vprn>log>log-id

```

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

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

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

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

If the AS number was previously changed, the BGP AS number inherits the new value.

- Special Cases**
- Service Admin State** — Bindings to an SDP within the service will be put into the out-of-service state when the service is shutdown. While the service is shutdown, all customer packets are dropped and counted as discards for billing and debugging purposes.  
A service is regarded as operational providing that one IP Interface SAP and one SDP is operational.
  - VPRN BGP and RIP** — This command disables the BGP or RIP instance on the given IP interface. Routes learned from a neighbor that is shutdown are immediately removed from the BGP or RIP database and RTM. If BGP or RIP is globally shutdown, then all RIP group and neighbor interfaces are shutdown operationally. If a BGP or RIP group is shutdown, all member neighbor interfaces are shutdown operationally. If a BGP or RIP neighbor is shutdown, just that neighbor interface is operationally shutdown.

## description

- Syntax** **description** *description-string*  
**no description**
- Context**
- ```

config>service>vprn
config>service>vprn>aarp-interface>spoke-sdp
config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
config>service>vprn>dhcp
config>service>vprn>dhcp>server>pool
config>service>vprn>if>dhcp
config>service>vprn>if>dhcp5
config>service>vprn>if>sap>ipsec-tunnel
config>service>vprn>l2tp
config>service>vprn>radius-proxy>server
config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor
config>service>vprn>ripng
config>service>vprn>ripng>group
config>service>vprn>ripng>group>neighbor
config>service>vprn>sub-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp
config>service>vprn>sub-if>grp-if>pppoe
config>service>vprn>sub-if>grp-if>sap>atm

```
- Description** This command creates a text description stored in the configuration file for a configuration context.
- The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

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

Default	no description
Parameters	<i>string</i> — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

description

Syntax	description <i>long-description-string</i> no description
Context	config>service>vprn>aarp-interface config>service>vprn>if config>service>vprn>if>sap config>service>vprn>nw-if config>service>vprn>red-if config>service>vprn>subscriber-interface config>service>vprn>sub-if>grp-if config>service>vprn>sub-if>grp-if>sap
Description	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the content in the configuration file. The no form of this command removes the string from the configuration.
Default	no description
Parameters	<i>string</i> — The description character string. Allowed values are any string up to 160 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

3.8.2.2 Global Commands

vprn

Syntax	vprn <i>service-id</i> [name <i>name</i>] [customer <i>customer-id</i>] [create] no vprn <i>service-id</i>

Context	config>service				
Description	This command creates or edits a Virtual Private Routed Network (VPRN) service instance. If the <i>service-id</i> does not exist, a context for the service is created. If the <i>service-id</i> 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 <i>customer-id</i> must be specified and associates the service with a customer. The <i>customer-id</i> must already exist having been created using the customer command in the service context. When a service is created with a customer association, it is not possible to edit the customer association. The service must be deleted and re-created with a new customer association. When a service is created, the use of the customer <i>customer-id</i> is optional to navigate into the service configuration context. Attempting to edit a service with the incorrect <i>customer-id</i> 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 a VPRN service ID belong to the same customer. The no form of the command deletes the VPRN service instance with the specified <i>service-id</i>. The service cannot be deleted until all the IP interfaces and all routing protocol configurations defined within the service ID have been shutdown and deleted.				
Parameters	<i>service-id</i> — 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 <i>service-id</i> must be the same number used for every 7750 SR on which this service is defined. Values <table><tr><td><i>service-id</i>:</td><td>1 to 2147483647</td></tr><tr><td><i>svc-name</i>:</td><td>64 characters maximum</td></tr></table> name <i>name</i> — Configures an optional service name identifier, up to 64 characters, to a given service. This service name can then be used in configuration references, display, and show commands throughout the system. A defined service name can help the service provider or administrator to identify and manage services within the SR OS platforms. To create a service, you must assign a service ID; however, after it is created, either the service ID or the service name can be used to identify and reference a service.	<i>service-id</i> :	1 to 2147483647	<i>svc-name</i> :	64 characters maximum
<i>service-id</i> :	1 to 2147483647				
<i>svc-name</i> :	64 characters maximum				

If a name is not specified at creation time, then SR OS assigns a string version of the *service-id* as the name.

Values *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 <i>ip-prefix/ip-prefix-length</i> [summary-only] [as-set] [aggregator <i>as-number:ip-address</i>] [black-hole [generate-icmp]] [community <i>comm-id</i> [<i>comm-id</i>] [local-preference <i>local-pref</i>]] [description <i>description</i>] aggregate <i>ip-prefix/ip-prefix-length</i> [summary-only] [as-set] [aggregator <i>as-number:ip-address</i>] [community <i>comm-id</i> [<i>comm-id</i>]] [indirect <i>ip-address</i>] [local-preference <i>local-pref</i>]] [description <i>description</i>] no aggregate <i>ip-prefix/ip-prefix-length</i>
Context	config>service>vprn
Description	This command creates an aggregate route. Use this command to automatically install an aggregate route in the routing table when there are one or more component routes. A component route is any route used for forwarding that is a more specific match of the aggregate. The use of aggregate routes can reduce the number of routes that need to be advertised to neighbor routers, leading to smaller routing table sizes. Overlapping aggregate routes may be configured; in this case a route becomes a component of only the one aggregate route with the longest prefix match. For example if one aggregate is configured as 10.0.0.0/16 and another as 10.0.0.0/24, then route 10.0.128/17 would be aggregated into 10.0.0.0/16, and route 10.0.0.128/25 would be aggregated into 10.0.0.0/24. If multiple entries are made with the same prefix and the same mask the previous entry is overwritten. A list of up to 12 BGP communities (any mix of standard, extended, and large communities) may be associated with an aggregate route. These communities can be matched in route policies and are automatically added to BGP routes that are created from the aggregate route. By default, aggregate routes are not installed in the forwarding table, however there are configuration options that allow an aggregate route to be installed with a black-hole next hop or with an indirect IP address as next hop.

Aggregate routes can be advertised via MP-BGP to other PEs within the network. Aggregate routes advertised using MP-BGP do not include aggregated BGP path attributes from the component routes which were used to activate the aggregate route. The aggregate route will be advertised with the minimal set of path attributes as if the aggregate was originated by the advertising routes. Export route policies should be used to control and modify the advertisement and path attributes of the aggregate routes.

The **no** form of the command removes the aggregate.

Default no aggregate

Parameters *ip-prefix* — The destination address of the aggregate route in dotted decimal notation.

Values

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 to 32
ipv6-prefix	x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
ipv6-prefix-length	0 to 128

the ipv6-prefix and ipv6-prefix-length apply only to the 7750 SR and 7950 XRS

the mask associated with the network address expressed as a mask length

Values: 0 to 32

summary-only — This optional parameter suppresses advertisement of more specific component routes for the aggregate.

To remove the **summary-only** option, enter the same aggregate command without the **summary-only** parameter.

as-set — This optional parameter is only applicable to BGP and creates an aggregate where the path advertised for this route will be an AS_SET consisting of all elements contained in all paths that are being summarized. Use this feature carefully as it can increase the amount of route churn due to best path changes.

aggregator as-number:ip-address — This optional parameter specifies the BGP aggregator path attribute to the aggregate route. When configuring the aggregator, a two-octet AS number used to form the aggregate route must be entered, followed by the IP address of the BGP system that created the aggregate route.

community — This configuration option associates a BGP community with the aggregate route. The community can be matched in route policies and is automatically added to BGP routes exported from the aggregate route.

comm-id — Specifies a BGP community value, up to 72 characters.

Values [as-num:comm-val | well-known-comm | ext-comm | large-comm]

where:

- *as-num* — 0 to 65535
- *comm-val* — 0 to 65535
- *well-known-comm* — **null** | **no-export** | **no-export-subconfed** | **no-advertise** | **llgr-stale** | **no-llgr** | **blackhole**
- *ext-comm* — the extended community, defined as one of the following:
 - {*target* | *origin*}:*ip-address*:*comm-val*
 - {*target* | *origin*}:*asnum*:*ext-comm-val*
 - {*target* | *origin*}:*ext-asnum*:*comm-val*
 - **bandwidth**:*asnum*:*val-in-mbps*
 - **ext:4300**:*ovstate*
 - **ext**:*value1*:*value2*
 - **color**:*co-bits*:*color-value*

where:

- *target* — route target
- *origin* — route origin
- *ip-address* — a.b.c.d
- *ext-comm-val* — 0 to 4294967295
- *ext-asnum* — 0 to 4294967295
- *val-in-mbps* — 0 to 16777215
- *ovstate* — 0, 1, or 2 (0 for valid, 1 for not found, 2 for invalid)
- *value1* — 0000 to FFFF
- *value2* — 0 to FFFFFFFFFF
- *co-bits* — 00, 01, 10 or 11
- *color-value* — 0 to 4294967295
- *large-comm* — *asn-or-ex:val-or-ex:val-or-ex*

black-hole — This optional parameter installs the aggregate route, when activated, in the FIB with a black-hole next-hop, where packets matching this route are discarded.

generate-icmp — This optional parameter keyword generates an ICMP.

indirect ip-address — This configuration option specifies that the aggregate route should be installed in the FIB with a next-hop taken from the route used to forward packets to ip-address.

Values

ipv4-prefix	a.b.c.d
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

The ipv6-prefix applies only to the 7750 SR and 7950 XRS.

local-preference *local-pref* — Specifies a BGP local-preference value with the aggregate route. The local-preference overrides the default local preference value of a BGP route originated by exporting the aggregate route.

Values 0 to 4294967295

description *description-text* — Specifies a text description stored in the configuration file for a configuration context.

allow-export-bgp-vpn

Syntax [no] **allow-export-bgp-vpn**

Context config>service>vpn

Description This command allows routes leaked from another local VPRN service to be re-exported by this VPRN in the form of new VPN-IP routes. The service label, route targets, and BGP next-hop of the re-advertised routes are based on the configuration and default values of the re-exporting VPRN.

When re-exporting leaked routes, the following restrictions apply.

- The **allow-export-bgp-vpn** command is not configurable in combination with any of the following commands: **carrier-carrier-vpn** (CSC), **label-mode next-hop** (LPN), **type {hub | spoke | subscriber-split-horizon}**, **redundant-interface**, and **export-inactive-bgp**.
- Re-exported routes always have the per-VRF label of the exporting VPRN; label-per-prefix advertisement is not supported.
- The best-external (inactive BGP) routes leaked by another VPRN cannot be re-exported by a VPRN configured with **allow-export-bgp-vpn**.



Caution: When a VPRN configured with **allow-export-bgp-vpn** advertises a leaked route, the **split-horizon** context is lost. A re-exported route can be easily advertised back to the sending peer unless this is blocked by BGP export policies. This can cause route flaps or other similar instability.

If the **no** form of the command is configured, leaked routes cannot be re-advertised as VPN-IP routes; they can only be re-advertised to PE-CE BGP peers of the VPRN.

Default no allow-export-bgp-vpn

auto-bind-tunnel

Syntax **auto-bind-tunnel**

Context	config>service>vprn
Description	This command enters the context to configure automatic binding of a VPRN service using tunnels to MP-BGP peers. The auto-bind-tunnel node is simply a context to configure the binding of VPRN routes to tunnels. The user must configure the resolution option to enable auto-bind resolution to tunnels in TTM. If the resolution option is explicitly set to disabled, the auto-binding to tunnel is removed. If resolution is set to any, any supported tunnel type in VPRN context will be selected following TTM preference. If one or more explicit tunnel types are specified using the resolution-filter option, then only these tunnel types will be selected again following the TTM preference. The user must set resolution to filter to activate the list of tunnel-types configured under resolution-filter. When an explicit SDP to a BGP next-hop is configured in a VPRN service (config>service>vprn>spoke-sdp), it overrides the auto-bind-tunnel selection for that BGP next-hop only. There is no support for reverting automatically to the auto-bind-tunnel selection if the explicit SDP goes down. The user must delete the explicit spoke-sdp in the VPRN service context to resume using the auto-bind-tunnel selection for the BGP next-hop.

ecmp

Syntax	ecmp <i>max-ecmp-routes</i> no ecmp
Context	config>service>vprn>auto-bind-tunnel
Description	This command configures the maximum number of tunnels that may be used as ECMP next-hops for the VPRN. This value overrides any values that have been configured using the config>service>vprn>ecmp command. The no form of the command removes the configured overriding value, and the value configured using the config>service>vprn>ecmp command will be used.
Default	no ecmp
Parameters	<i>max-ecmp-routes</i> — Specifies the maximum number of tunnels that may be used as ECMP next-hops for the VPRN. Values 0 to 32

enforce-strict-tunnel-tagging

Syntax	[no] enforce-strict-tunnel-tagging

Context	config>service>vprn>auto-bind-tunnel
Description	This command forces the system to only consider LSPs marked with an admin-tag for next hop resolution. Untagged LSPs are not be considered. The no form of this command reverts to the default behavior. While tagged RSVP and SR-TE LSPs are considered first, the system can fall back to using untagged LSP of other types and not exclude them as per the auto-bind-tunnel configuration.
Default	no enforce-strict-tunnel-tagging

resolution

Syntax	resolution {any filter disabled}
Context	config>service>vprn>auto-bind-tunnel
Description	This command configures the resolution mode in the automatic binding of a VPRN service to tunnels to MP-BGP peers.
Parameters	<i>any</i> — Enables the binding to any supported tunnel type in VPRN context following TTM preference. <i>filter</i> — Enables the binding to the subset of tunnel types configured under resolution-filter. <i>disabled</i> — Disables the automatic binding of a VPRN service to tunnels to MP-BGP peers.

resolution-filter

Syntax	resolution-filter
Context	config>service>vprn>auto-bind-tunnel
Description	This command configures the subset of tunnel types that can be used in the resolution of VPRN prefixes within the automatic binding of VPRN service to tunnels to MP-BGP peers. The following tunnel types are supported in a VPRN context in order of preference: RSVP, SR-TE, GRE, RIB-API, LDP, SR-ISIS, SR-OSPF, SR-policy, BGP, and UDP.

bgp

Syntax	[no] bgp
Context	config>service>vprn>auto-bind-tunnel>resolution-filter
Description	Selects the BGP tunnel type.

The **bgp** value instructs BGP IP-VPN to search for a BGP LSP to the address of the BGP next hop. If the user does not enable the BGP tunnel type, inter-area or inter-as prefixes will not be resolved.

gre

Syntax	[no] gre
Context	config>service>vprn>auto-bind-tunnel>resolution-filter
Description	Selects the GRE tunnel type. The gre encapsulation of the MPLS service packet uses the base 4-byte header as per RFC 2890. The optional fields Checksum (plus Reserved field), Key, and Sequence Number are not inserted.

ldp

Syntax	[no] ldp
Context	config>service>vprn>auto-bind-tunnel>resolution-filter
Description	Selects the LDP tunnel type. The ldp value instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next hop.

rib-api

Syntax	[no] rib-api
Context	config>service>vprn>auto-bind-tunnel>resolution-filter
Description	Selects the RIB-API tunnel type. The rib-api value allows tunnels programmed using the RibApi gRPC service for use in resolving the next hops of VPN-IPv4 and VPN-IPv6 routes imported into the VPRN service.

rsvp

Syntax	[no] rsvp
Context	config>service>vprn>auto-bind-tunnel>resolution-filter
Description	Selects the RSVP-TE tunnel type.

The **rsvp** value instructs BGP to search for the best metric RSVP LSP to the address of the BGP next hop. This address can correspond to the system interface or to another loopback interface 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**.

sr-isis

Syntax	[no] sr-isis
Context	config>service>vprn>auto-bind-tunnel>resolution-filter
Description	Selects the Segment Routing (SR) tunnel type programed by an IS-IS instance in TTM. When the sr-isis (or sr-ospf) value is enabled, an SR tunnel to the BGP next hop is selected in the TTM from the lowest-numbered IS-IS (OSPF) instance.

sr-ospf

Syntax	[no] sr-ospf
Context	config>service>vprn>auto-bind-tunnel>resolution-filter
Description	Selects the SR-OSPF tunnel type. When the sr-ospf (sr-isis) value is enabled, a SR tunnel to the BGP next hop is selected in the TTM from the lowest numbered IS-IS (OSPF) instance.

sr-policy

Syntax	[no] sr-policy
Context	config>service>vprn>auto-bind-tunnel>resolution-filter
Description	Selects the SR policy tunnel type. The sr-policy value instructs BGP to search for an SR policy with a non-null endpoint and color value that matches the BGP next hop and color extended community value, respectively, of the VPN-IP route.

sr-te

Syntax	[no] sr-te
Context	config>service>vprn>auto-bind-tunnel>resolution-filter

Description Selects the SR-TE tunnel type.

The **sr-te** value instructs the system to search for the best metric SR-TE LSP to the address of the BGP next hop. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple SR-TE LSPs with the same lowest metric, BGP selects the LSP with the lowest **tunnel-id**.

udp

Syntax [no] **udp**

Context config>service>vprn>auto-bind-tunnel>resolution-filter

Description Selects the UDP tunnel type.

The **udp** value instructs BGP EVPN to search for a UDP LSP to the address of the BGP next hop.

weighted-ecmp

Syntax [no] **weighted-ecmp**

Context config>service>vprn>auto-bind-tunnel

Description This command enables weighted ECMP for packets using tunnels that a VPRN automatically binds to. When weighted ECMP is enabled, packets are sprayed across LSPs in the ECMP according to the outcome of the hash algorithm and the configured **load-balancing-weight** of each LSP.

The **no** form of the command disables weighted ECMP for next hop tunnel selection.

Default no weighted-ecmp

autonomous-system

Syntax **autonomous-system** *as-number*
no autonomous-system

Context config>service>vprn

Description This command defines the autonomous system (AS) to be used by this VPN routing/forwarding (VRF). This command defines the autonomous system to be used by this VPN routing

The **no** form of the command removes the defined AS from this VPRN context.

Default no autonomous-system

Parameters	<i>as-number</i> — Specifies the AS number for the VPRN service.
Values	1 to 4294967295

backup-path

Syntax	[no] backup-path [ipv4] [ipv6] [label-ipv4] [label-ipv6]
Context	config>router config>service>vprn config>service>vprn>bgp
Description	This command enables the computation and use of a backup path for IPv4 and/or IPv6 BGP-learned prefixes belonging to the base router or a particular VPRN. Multiple paths must be received for a prefix in order to take advantage of this feature. When a prefix has a backup path and its primary path(s) fail the affected traffic is rapidly diverted to the backup path without waiting for control plane re-convergence to occur. When many prefixes share the same primary path(s), and in some cases also the same backup path, the time to failover traffic to the backup path is independent of the number of prefixes. By default, IPv4 and IPv6 prefixes do not have a backup path installed in the IOM.
Default	no backup-path
Parameters	ipv4 — Enables the use of a backup path for BGP-learned unlabeled IPv4 prefixes. ipv6 — Enables the use of a backup path for BGP-learned unlabeled IPv6 prefixes. label-ipv4 — Enables the use of a backup path for BGP-learned labeled-IPv4 prefixes. label-ipv6 — Enables the use of a backup path for BGP-learned labeled-IPv6 prefixes. label-ipv6 is not supported within the config>service>vprn context.

carrier-carrier-vpn

Syntax	[no] carrier-carrier-vpn
Context	config>service>vprn
Description	This command configures a VPRN service to support a Carrier Supporting Carrier model. It should be configured on a network provider's CSC-PE device. This command cannot be applied to a VPRN unless it has no SAP or spoke-SDP interfaces. Once this command has been entered one or more MPLS-capable CSC interfaces can be created in the VPRN. The no form of the command removes the Carrier Supporting Carrier capability from a VPRN.
Default	no carrier-carrier-vpn

confederation

Syntax	confederation <i>confed-as-num</i> members <i>as-number</i> [<i>as-number</i>] no confederation <i>confed-as-num</i> members <i>as-number</i> [<i>as-number</i>] no confederation
Context	config>service>vprn
Description	This command configures the VPRN BGP instance to participate in a BGP confederation. BGP confederations can be used to reduce the number of IBGP sessions required within an AS. When a VPRN BGP instance is part of a confederation, it can form confederation-EBGP sessions with CE router peers in a different sub-autonomous systems of the same confederation as well as regular EBGP sessions with CE router peers outside the confederation. A VPRN BGP instance that is part of a confederation cannot import or export its routes to the base router instance (as VPN-IP routes). The no form of the command deletes the specified member AS from the confederation. When members are not specified in the no statement, the entire list is removed and confederations is disabled. When the last member of the list is removed, confederations is disabled.
Default	no confederation
Parameters	<i>confed-as-num</i> — The confederation AS number defined as a decimal value. Values 1 to 4294967295 members <i>as-number</i> — The AS number(s) that are members of the confederation, each expressed as a decimal integer. Configure up to 15 members per confed-as-num. Values 1 to 4294967295

dns

Syntax	[no] dns
Context	config>service>vprn
Description	This command enters the context to configure domain name servers. The no form of the command disables DNS for this service.

ipv4-source-address

Syntax	ipv4-source-address <i>ipv4-address</i> no ipv4-source-address
Context	config>service>vprn>dns

Description	This command configures the IPv4 address of the default secondary DNS server for the subscribers using this interface. Subscribers that cannot obtain an IPv4 DNS server address by other means, can use this for DNS name resolution. The <code>ipv4-address</code> value can only be set to a nonzero value if the value of VPRN type is set to subscriber-split-horizon. The no form of the command reverts to the default.
Parameters	<i>ipv4-address</i> — Specifies the IPv4 address of the default secondary DNS server. Values <i>ipv4-address</i> - a.b.c.d

ipv6-source-address

Syntax	ipv6-source-address <i>ipv6-address</i> no ipv6-source-address
Context	config>service>vprn>dns
Description	This command configures the IPv6 address of the default secondary DNS server for the subscribers using this interface. Subscribers that cannot obtain an IPv6 DNS server address by other means, can use this for DNS name resolution. The <code>ipv6-address</code> value can only be set to a nonzero value if the value of VPRN type is set to subscriber-split-horizon. The no form of the command reverts to the default.
Parameters	<i>ipv6-address</i> — Specifies the IPv6 address of the default secondary DNS server. Values <i>ipv6-address</i> - a.b.c.d

primary-dns

Syntax	primary-dns <i>ip-address</i> no primary-dns
Context	config>service>vprn>dns
Description	This command configures the primary DNS server used for DNS name resolution. DNS name resolution can be used when executing ping, traceroute, and service-ping, and also when defining file URLs. DNS name resolution is not supported when DNS names are embedded in configuration files. The no form of the command removes the primary DNS server from the configuration.
Default	no primary-dns — No primary DNS server is configured.

Parameters *ip-address* — The IP or IPv6 address of the primary DNS server.

Values

ipv4-address -a.b.c.d

ipv6-address: x:x:x:x:x:x:x[-interface]

x:x:x:x:x:d.d.d.d[-interface]

x: [0..FFFF]H

d: [0..255]D

interface - 32 characters max, for link local addresses.

secondary-dns

Syntax **secondary-dns** *ip-address*
 no secondary-dns

Context config>service>vprn>dns

Description This command configures the secondary DNS server for DNS name resolution. The secondary DNS server is used only if the primary DNS server does not respond.

DNS name resolution can be used when executing ping, traceroute, and service-ping, and also when defining file URLs. DNS name resolution is not supported when DNS names are embedded in configuration files.

The **no** form of the command removes the secondary DNS server from the configuration.

Default no secondary-dns — No secondary DNS server is configured.

Parameters *ip-address* — The IP or IPv6 address of the secondary DNS server.

Values

ipv4-address -a.b.c.d

ipv6-address: x:x:x:x:x:x:x[-interface]

x:x:x:x:x:d.d.d.d[-interface]

x: [0 to FFFF]H

d: [0 to 255]D

interface - 32 characters max, for link local addresses.

tertiary-dns

Syntax **tertiary-dns** *ip-address*
 no tertiary-dns

Context config>service>vprn>dns

Description	This command configures the tertiary DNS server for DNS name resolution. The tertiary DNS server is used only if the primary DNS server and the secondary DNS server do not respond. DNS name resolution can be used when executing ping, traceroute, and service-ping, and also when defining file URLs. DNS name resolution is not supported when DNS names are embedded in configuration files. The no form of the command removes the tertiary DNS server from the configuration.
Default	no tertiary-dns — No tertiary DNS server is configured.
Parameters	<i>ip-address</i> — The IP or IPv6 address of the tertiary DNS server. Values ipv4-address -a.b.c.d ipv6-address: x:x:x:x:x:x:x[-interface] x:x:x:x:x:x:d.d.d.d[-interface] x: [0 to FFFF]H d: [0 to 255]D interface - 32 characters max, for link local addresses.

ecmp

Syntax	ecmp <i>max-ecmp-routes</i> no ecmp
Context	config>service>vprn
Description	This command enables equal-cost multipath (ECMP) and configures the number of routes for path sharing. For example, the value of 2 means that 2 equal cost routes will be used for cost sharing. ECMP groups form when the system routes to the same destination with equal cost values. Routing table entries can be entered manually (as static routes), or they can be formed when neighbors are discovered and routing table information is exchanged by routing protocols. The system can balance traffic across the groups with equal costs. ECMP can only be used for routes learned with the same preference and same protocol. See the discussion on preferences in the application6 command. When more ECMP routes are available at the best preference than configured by the max-ecmp-routes parameter, then the lowest next-hop IP address algorithm is used to select the number of routes configured. The no form of the command disables ECMP path sharing. If ECMP is disabled and multiple routes are available at the best preference and equal cost, the newly updated route is used.
Default	no ecmp

Parameters	<i>max-ecmp-routes</i> — Specifies the maximum number of routes for path sharing.
Values	0 to 32

export-grt

Syntax	export-grt <i>plcy-or-long-expr</i> [<i>plcy-or-expr</i> [<i>plcy-or-expr</i> ... (up to 4 max)]] no export-grt
Context	config>service>vprn
Description	This command is used to specify route policies that control how routes are exported from the local VRF route table to the base router (GRT) route table. The leaked routes show as protocol VPN-Leak in the GRT and allow traffic to ingress on a GRT interface and egress on a VPRN interface. The export-grt command can reference up to 5 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine final action to accept or reject the route. Only one of the 5 objects referenced by the export-grt command can be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 5 objects has a maximum length of 255 characters while the remaining 4 objects have a maximum length of 64 characters each. When multiple export-grt commands are issued, the last command entered overrides the previous command. The no form of the command removes all route policy names from the export-grt list.
Default	no export-grt
Parameters	<i>plcy-or-long-expr</i> — The route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long). <i>plcy-or-expr</i> — The route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long).

export-inactive-bgp

Syntax	export-inactive-bgp no export-inactive-bgp
Context	config>service>vprn
Description	This command allows the best BGP route learned by a VPRN to be exported as a VPN-IP route even when that BGP route is inactive in the route table due to the presence of a preferred BGP-VPN route from another PE. In order for the BGP route to be exported, it must be accepted by the VRF export policy.

This “best-external” type of route advertisement is useful in active/standby multi-homing scenarios because it can ensure that all PEs have knowledge of the backup path provided by the standby PE.

By default, an inactive BGP route cannot be exported from a VPRN.

Default no export-inactive-bgp

fib-priority

Syntax **fib-priority** {**high** | **standard**}

Context config>service>vprn

Description This command specifies the FIB priority for VPRN BGP routes.

Parameters **high** — Specifies high FIB priority for VPRN.
standard — Specifies standard FIB priority for VPRN.

flowspec

Syntax **flowspec**

Context config>service>vprn

Description This command enters the context to configure flowspec-related parameters for the specified routing instance.

ip-filter-max-size

Syntax **ip-filter-max-size** {*value* | **default**}

Context config>service>vprn>flowspec

Description This command configures the maximum number of flowspec routes or rules that can be embedded into an ingress IP filter policy for a specified routing instance. Flowspec filter entries embedded in a filter policy in this routing instance will use filter entries from the range between the embedding offset and “offset + ip-filter-max-size – 1”.

The sum of the **ip-filter-max-size** *value* parameter and the highest offset in any IPv4 filter that embeds IPv4 flowspec rules from this routing instance (excluding filters that embed at offset 262143) must not exceed 262143.

The **ip-filter-max-size** configuration can be adjusted up or down at any time. If the number of IPv4 flowspec rules that are currently installed is M , and the new limit is N , where $N < M$, then the last set of rules from N to M (by flowspec order) are immediately removed, but are retained in the BGP RIB. If the limit is increased, new rules are programmed only as they are received again in new BGP updates.

Default	ip-filter-max-size default
Parameters	<i>value</i> — The maximum number of flowspec routes or rules that can be embedded into an ingress IP filter policy.
Values	0 to 262143
default	— Configures the maximum size as 512.

ipv6-filter-max-size

Syntax	ipv6-filter-max-size { <i>value</i> default }
Context	config>service>vprn>flowspec
Description	This command configures the maximum number of IPv6 flowspec routes or rules that can be embedded into an ingress IPv6 filter policy for a specified routing instance. Flowspec filter entries embedded in a filter policy in this routing instance will use filter entries from the range between the embedding offset and “offset + ip-filter-max-size – 1”. The sum of the ip-filter-max-size <i>value</i> parameter and the highest offset in any IPv6 filter that embeds IPv6 flowspec rules from this routing instance (excluding filters that embed at offset 262143) must not exceed 262143. The ip-filter-max-size configuration can be adjusted up or down at any time. If the number of IPv6 flowspec rules that are currently installed is M, and the new limit is N, where $N < M$, then the last set of rules from N to M (by flowspec order) are immediately removed, but are retained in the BGP RIB. If the limit is increased, new rules are programmed only as they are received again in new BGP updates.
Default	ipv6-filter-max-size default
Parameters	<i>value</i> — The maximum number of flowspec routes or rules that can be embedded into an ingress IP filter policy.
Values	0 to 262143
default	— Configures the maximum size as 512.

grt-lookup

Syntax	grt-lookup
Context	config>service>vprn

Description This command provides the context under which all Global Route Table (GRT) leaking commands are configured. If all the supporting commands in the context are removed, this command will also be removed.

enable-grt

Syntax [no] enable-grt

Context config>service>vprn>grt-lookup

Description This command enables the functions required for looking up routes in the Global Route Table (GRT) when the lookup in the local VRF fails. If this command is enabled without the use of a **static-route** option (as subcommand to this parent), a lookup in the local VRF is preferred over the GRT. When the local VRF returns no route table lookup matches, the result from the GRT is preferred.

The **no** form of this command disables the lookup in the GRT when the lookup in the local VRF fails.

Default no enable-grt

allow-local-management

Syntax [no] allow-local-management

Context config>service>vprn>grt-lookup>enable-grt

Description Enables the support of specific management protocols over VPRN interfaces that terminate on Base routing context IPv4 and IPv6 interface addresses, including Base loopback and system addresses. Global Routing Table (GRT) leaking is used to enable visibility/access of the Base interface addresses in the VPRN. The supported protocols are Telnet, FTP, SNMP, and SSH (including applications that ride over SSH such as SCP and SFTP) and TACAS+.

Ping and traceroute responses from the Base router interfaces are supported and are not configurable.

The **allow-local-management** command does not control the support for management protocols terminating on VPRN interfaces directly. See the **config>service>vprn>snmp>access** CLI command for SNMP support on VPRN interface addresses.

export-grt

Syntax **export-grt** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]
no export-grt

Context	config>service>vprn>grt-lookup
Description	This command uses route policy to determine which routes are exported from the VRF to the GRT along with all the forwarding information. These entries will be marked as BGP-VPN routes in the GRT. Routes must be in the GRT in order for proper routing to occur from the GRT to the VRF.
Default	no export-grt
Parameters	<i>plcy-or-long-expr</i> — The route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long). <i>plcy-or-expr</i> — The route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long). Up to 4 policy names or logical expressions can be specified in a single statement.

export-limit

Syntax	export-limit <i>num-routes</i> no export-limit
Context	config>service>vprn>grt-lookup config>service>vprn>ospf config>service>vprn>ospf3 config>service>vprn>rip
Description	This command provides the ability to limit the total number of routes exported from the VRF to the GRT. The value zero (0) provides an override that disables the maximum limit. Setting this value to zero (0) will not limit the number of routes exported from the VRF to the GRT. Configuring a range of one (1) to 1000 will limit the number of routes to the specified value. The no form of the command sets the export-limit to a default of five (5).
Default	export-limit 5
Parameters	<i>num-routes</i> — Specifies the maximum number of routes that can be exported. Values 0 to 1000

export-v6-limit

Syntax	export-v6-limit <i>num-routes</i> no export-v6-limit
Context	config>service>vprn>grt-lookup

Description	The export-limit range provides the ability to limit the total number of IPv6 routes exported from the VPRN to the GRT. The value "0" provides an override that disables the maximum limit. Setting this value to "0" will not limit the number of routes exported from the VPRN to the GRT. Configuring a range of 1-1000 will limit the number of routes to the specified value. The no form of the command sets the export-limit to a default of 5.
Default	export-v6-limit 5
Parameters	<i>num-routes</i> — Specifies maximum number of routes that can be exported. Values 0 to 1000

Description This command specifies a GSMP name. A GSMP group name is unique only within the scope of the service in which it is defined.

Parameters *name* — Specifies the group name up to 32 characters in length.

ancp

Syntax **ancp**

Context config>service>vprn>gsmp>group

Description This command configures ANCP parameters for this GSMP group.

dynamic-topology-discover

Syntax [**no**] **dynamic-topology-discover**

Context config>service>vprn>gsmp>group>ancp

Description This command enables the ANCP dynamic topology discovery capability.

The **no** form of this command disables the feature.

oam

Syntax [**no**] **oam**

Context config>service>vprn>gsmp>group>ancp

Description This command specifies whether or not the GSMP ANCP OAM capability should be negotiated at startup of the GSMP connection.

The **no** form of this command disables the feature.

hold-multiplier

Syntax **hold-multiplier** *multiplier*
no hold-multiplier

Context config>service>vprn>gsmp>group

Description This command configures the hold-multiplier for the GSMP connections in this group.

Parameters *multiplier* — Specifies the GSMP hold multiplier value.

Values 1 to 100

idle-filter

Syntax	idle-filter no idle-filter
Context	config>service>vprn>gsmp
Description	This command when applied will filter out new subscriber's ANCP messages from subscriber with "DSL-line-state" IDLE.
Default	no idle-filter

keepalive

Syntax	keepalive <i>seconds</i> no keepalive
Context	config>service>vprn>gsmp>group
Description	This command configures keepalive values for the GSMP connections in this group.
Parameters	<i>seconds</i> — Specifies the GSMP keepalive timer value in seconds. Values 1 to 25

neighbor

Syntax	[no] neighbor <i>ip-address</i>
Context	config>service>vprn>gsmp>group
Description	This command adds or removes a neighbor in this group.
Parameters	<i>ip-address</i> — Specifies the IP address in dotted decimal notation.

local-address

Syntax	local-address <i>ip-address</i> no local-address
Context	config>service>vprn>gsmp>group>neighbor
Description	This command configures the source ip-address used in the connection towards the neighbor.
Parameters	<i>ip-address</i> — Specifies the IP address in dotted decimal notation.

priority-marking

Syntax	priority-marking dscp <i>dscp-name</i> priority-marking prec <i>ip-prec-value</i> no priority-marking
Context	config>service>vprn>gsmp>group>neighbor
Description	This command configures the type of priority marking to be used.
Parameters	<i>dscp dscp-name</i> — Specifies the DSCP code-point to be used. Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63 <i>prec ip-prec-value</i> — Specifies the precedence value to be used. Values 0 to 7

persistency-database

Syntax	[no] persistency-database
Context	config>service>vprn>gsmp
Description	This command enables the system to store DSL line information in memory. If the GSMP connection terminates, the DSL line information will remain in memory and accessible for Radius authentication and accounting.
Default	no persistency-database

3.8.2.3 Router L2TP Commands

l2tp

Syntax	[no] l2tp
Context	config>service>vprn
Description	This command enters the context to configure L2TP parameters. L2TP extends the PPP model by allowing Layer 2 and PPP endpoints to reside on different devices interconnected by a packet-switched network.

avp-hiding

Syntax	avp-hiding <i>sensitive</i> <i>always</i> no avp-hiding
Context	config>service>vprn>l2tp
Description	This command configures Attribute Value Pair (AVP) hiding. This capability can be used to avoid the passing of sensitive data, such as user passwords, as clear text in an AVP. The no form of the command returns the value to never allow AVP hiding.
Default	no avp-hiding
Parameters	<i>avp-hiding</i> — Specifies the method to be used for the authentication of the tunnels in this L2TP group. Values <i>sensitive</i> — AVP hiding is used only for sensitive information (such as username/password). <i>always</i> — AVP hiding is always used.

calling-number-format

Syntax	calling-number-format <i>ascii-spec</i> no calling-number-format
Context	config>service>vprn>l2tp
Description	This command what string to put in the Calling Number AVP, for L2TP control messages related to a session in this L2TP protocol instance.
Parameters	<i>ascii-spec</i> — Specifies the L2TP calling number AVP. Values <i>ascii-spec</i> char-specification <i>ascii-spec</i> char-specification <i>ascii-char</i> <i>char-origin</i> <i>ascii-char</i> a printable ASCII character <i>char-origin</i> %origin <i>origin</i> S c r s l S - system name, the value of TIMETRA-CHASSIS-MIB::tmnxChassisName c - Agent Circuit Id r - Agent Remote Id s - SAP ID, formatted as a character string l - Logical Line ID

challenge

Syntax	challenge { always never } no challenge
Context	config>service>vprn>l2tp
Description	This command configures the use of challenge-response authentication. The no form of the command reverts to the default never value.
Default	no challenge
Parameters	always — Specifies that challenge-response authentication is always used. never — Specifies that challenge-response authentication is never used.

destruct-timeout

Syntax	destruct-timeout <i>destruct-timeout</i> no destruct-timeout
Context	config>service>vprn>l2tp
Description	This command configures the period of time that the data of a disconnected tunnel will persist before being removed. The no form of the command removes the value from the configuration.
Default	no destruct-timeout
Parameters	<i>destruct-timeout</i> — Specifies the automatic removal of dynamic L2TP sessions, in seconds, that are no longer active. Default no destruct-timeout Values 60 to 86400

exclude-avps

Syntax	exclude-avps <i>calling-number</i> no exclude-avps
Context	config>service>vprn>l2tp
Description	This command configures the L2TP AVPs to exclude.

ipcp-subnet-negotiation

Syntax	[no] ipcp-subnet-negotiation
Context	config>service>vprn>l2tp>group>ppp config>service>vprn>l2tp>group>tunnel>ppp
Description	Enables IPCP negotiation for PPPoE hosts. If not enabled (default setting), the current behavior will apply even if subnet is allocated to the host. Enables IPCP negotiation for PPPoE hosts. If not enabled (default setting), the current behavior will apply even if subnet is allocated in the host.

peer-address-change-policy

Syntax	peer-address-change-policy {accept ignore reject}
Context	config>service>vprn>l2tp
Description	This command configures the reaction to a change of tunnel peer address in this router.

receive-window-size

Syntax	receive-window-size <i>window-size</i> no receive-window-size
Context	config>service>vprn>l2tp
Description	This command configures the L2TP receive window size.
Parameters	<i>window-size</i> — Specifies the window size. Values 4 to 1024

rtm-debounce-time

Syntax	rtm-debounce-time <i>debounce-time</i> no rtm-debounce-time
Context	config>service>vprn>l2tp
Description	This command configures the amount of time, in milliseconds, that the system will wait before declaring an L2TP tunnel down when the remote endpoint IP address cannot be resolved to an active IP route in the local routing table. The default behavior is for the L2TP tunnel to not be declared down based on the remote endpoint IP address reachability.

The **no** form of this command returns the **rtm-debounce-time** to the default value of zero.

Default	no rtm-debounce-time
Parameters	<i>debounce-time</i> — Specifies the amount of time, in milliseconds, that the system will wait before declaring the associated L2TP tunnel as down.
Values	0 to 5000

group

Syntax	group <i>tunnel-group-name</i> [create] no group <i>tunnel-group-name</i>
Context	config>service>vprn>l2tp
Description	This command configures an L2TP tunnel group.
Parameters	<i>tunnel-group-name</i> — Specifies a name string to identify a L2TP group up to 63 characters in length. create — This keyword is mandatory when creating a tunnel group name. The create keyword requirement can be enabled/disabled in the environment>create context.

session-limit

Syntax	session-limit <i>session-limit</i> session-limit unlimited no session-limit
Context	config>service>vprn>l2tp
Description	This command configures the L2TP session limit for the router. L2TP is connection-oriented. The L2TP Network Server (LNS) and LAC maintain state for each call that is initiated or answered by an LAC. An L2TP session is created between the LAC and LNS when an end-to-end PPP connection is established between a remote system and the LNS. Datagrams related to the PPP connection are sent over the tunnel between the LAC and LNS. There is a one to one relationship between established L2TP sessions and their associated calls.
Default	no session-limit
Parameters	<i>session-limit</i> — Specifies the number of sessions allowed. Values 1 to 131071 unlimited — Specifies the use of the maximum available number of sessions allowed.

avp-hiding

Syntax	avp-hiding <i>sensitive</i> <i>always</i> no avp-hiding
Context	config>service>vprn>l2tp>group
Description	This command configures Attribute Value Pair (AVP) hiding. This capability can be used to avoid the passing of sensitive data, such as user passwords, as clear text in an AVP. The no form of the command returns the value to never allow AVP hiding.
Default	no avp-hiding
Parameters	<i>avp-hiding</i> — Specifies the method to be used for the authentication of the tunnels in this L2TP group. Values <i>sensitive</i> — AVP hiding is used only for sensitive information (such as username/password). <i>always</i> — AVP hiding is always used.

challenge

Syntax	challenge <i>always</i> no challenge
Context	config>service>vprn>l2tp>group
Description	This command configures the use of challenge-response authentication. The no form of the command reverts to the default never value.
Default	no challenge
Parameters	<i>always</i> — Specifies when challenge-response is to be used for the authentication of the tunnels in this L2TP group. Values <i>always</i>

destruct-timeout

Syntax	destruct-timeout <i>destruct-timeout</i> no destruct-timeout
Context	config>service>vprn>l2tp>group config>service>vprn>l2tp>group>tunnel
Description	This command configures the period of time that the data of a disconnected tunnel will persist before being removed.

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

Default	no destruct-timeout
Parameters	<i>destruct-timeout</i> — Specifies the automatic removal of dynamic L2TP sessions, in seconds, that are no longer active.
Values	60 to 86400

hello-interval

Syntax	hello-interval <i>hello-interval</i> no hello-interval
Context	config>service>vprn>l2tp>group
Description	This command configures the time interval between two consecutive tunnel Hello messages. The Hello message is an L2TP control message sent by either peer of a LAC-LNS control connection. This control message is used as a keepalive for the tunnel. The no form of the command removes the interval from the configuration.
Default	no hello-interval
Parameters	<i>hello-interval</i> — Specifies the time interval, in seconds, between two consecutive tunnel Hello messages. Values 60 to 3600 Default 60

idle-timeout

Syntax	idle-timeout <i>idle-timeout</i> no idle-timeout
Context	config>service>vprn>l2tp>group
Description	This command configures the period of time that an established tunnel with no active sessions will persist before being disconnected. Enter the no form of the command to maintain a persistent tunnel. The no form of the command removes the idle timeout from the configuration.
Default	no idle-timeout
Parameters	<i>idle-timeout</i> — Specifies the idle timeout value, in seconds until the group is removed. Values 0 to 3600

l2tpv3

Syntax	l2tpv3
Context	config>service>vprn>l2tp config>service>vprn>l2tp>group
Description	This command enters the context to configure L2TPv3 parameters.

cookie-length

Syntax	cookie-length {4 8 default} no cookie-length
Context	config>service>vprn>l2tp>l2tpv3 config>service>vprn>l2tp>group>l2tpv3
Description	This command configures the length of the optional cookie field. The no form of the command returns the cookie-length to a default of none .
Default	no cookie-length
Parameters	4 — Specifies the cookie length as 4 bytes. 8 — Specifies the cookie length as 8 bytes. default — When specified within the config>service>vprn>l2tp>group>l2tpv3 context, this is referencing to the cookie-length configuration within the config>service>vprn>l2tp>l2tpv3 context.

digest-type

Syntax	digest-type {default none md5 sha1} no digest-type
Context	config>service>vprn>l2tp>l2tpv3 config>service>vprn>l2tp>group>l2tpv3
Description	This command configures the hashing algorithm used to calculate the message digest. The no form of the command returns the digest-type to none .
Default	no digest-type
Parameters	none — Specifies that no digest should be used. md5 — Specifies that the MD5 algorithm should be used. sha1 — Specifies that the SHA1 algorithm should be used.

default — When specified within the **config>service>vprn>l2tp>group>l2tpv3** context, this is referencing to the **digest-type** configuration within the **config>service>vprn>l2tp>l2tpv3** context.

nonce-length

Syntax	nonce-length { <i>length</i> default } no nonce-length
Context	config>service>vprn>l2tp>l2tpv3 config>service>vprn>l2tp>group>l2tpv3
Description	This command configures the length for the local L2TPv3 nonce (random number) value used in the Nonce AVP. The no form of the command returns the nonce-length to a default of none .
Default	no nonce-length
Parameters	<i>length</i> — Specifies the length of the Nonce AVP value. Values 16 to 64 default — When specified within the config>service>vprn>l2tp>group>l2tpv3 context, this is referencing to the nonce-length configuration within the config>service>vprn>l2tp>l2tpv3 context.

private-tcp-mss-adjust

Syntax	private-tcp-mss-adjust <i>octets</i> no private-tcp-mss-adjust
Context	config>service>vprn>l2tp>l2tpv3
Description	This command enables TCP MSS adjust for L2TPv3 tunnels on the private side of the service level. When this command is configured, the system updates the TCP MSS option value of the received TCP SYN packet on the private side. Note that this command can be overridden by the corresponding configuration on the group or tunnel level. The no form of this command disables TCP MSS adjust on the private side.
Default	no private-tcp-mcc-adjust
Parameters	<i>octets</i> — Specifies the new TCP MSS value in octets. Values 512 to 9000

public-tcp-mss-adjust

Syntax	public-tcp-mss-adjust <i>octets</i> no public-tcp-mss-adjust
Context	config>service>vprn>l2tp>l2tpv3
Description	This command enables TCP MSS adjust for L2TPv3 tunnels on the public side on the service level. When the command is configured, the system updates the TCP MSS option value of the received TCP SYN packet on the public side that is encapsulated in the L2TPv3 tunnel. Note that this command can be overridden by the corresponding configuration on the group or tunnel level. The no form of this command disables TCP MSS adjust on the public side.
Default	no public-tcp-mss-adjust
Parameters	<i>octets</i> — Specifies the new TCP MSS value in octets
Values	512 to 9000

rem-router-id

Syntax	rem-router-id <i>ip-addr</i> no rem-router-id
Context	config>service>vprn>l2tp>group>l2tpv3
Description	This command configures the IP address that should be used within the Remote Router-ID AVP. The no form of this command removes the configured IP address.
Default	no rem-router-id
Parameters	<i>ip-addr</i> — Specifies an IP address to be used within the Remote Router-ID AVP.

pw-cap-list

Syntax	pw-cap-list { ethernet ethernet-vlan } no pw-cap-list
Context	config>service>vprn>l2tp>group>l2tpv3
Description	This command configures the allowable pseudowire capability list that is advertised to the far end. An empty list results in both pseudowire capabilities being advertised.

The **no** form of this command removes the list and advertises both pseudowire capabilities to the far end.

- Default** no pw-cap-list
- Parameters** **ethernet** — Specifies that the Ethernet pseudo-wire type is advertised.
ethernet-vlan — Specifies that the Ethernet-VLAN pseudo-wire type is advertised.

track-password-change

- Syntax** [no] track-password-change
- Context** config>service>vprn>l2tp>group>l2tpv3
- Description** This command enables tracking of password changes, allowing password tunnel passwords to be changed without bringing down active tunnels or sessions. This is only supported with L2TPv3.
- The **no** form of the command disables password change tracking.
- Default** no track-password-change

transport-type

- Syntax** transport-type ip
no transport-type
- Context** config>service>vprn>l2tp>l2tpv3
- Description** This command configures the transport type to be used to carry the L2TPv3 tunnel. Currently, only IP transport is supported.
- The **no** form of this command returns the **transport-type** to the default value.
- Default** no transport-type
- Parameters** **ip** — Specifies that IP should be used as the transport type for the L2TPv3 tunnel.

lns-group

- Syntax** lns-group lns-group-id
no lns-group
- Context** config>service>vprn>l2tp>group
- Description** This command configures the ISA LNS group.

Parameters	<i>lms-group-id</i> — Specifies the LNS group ID.
Values	1 to 4

local-address

Syntax	local-address <i>ip-address</i> no local-address
Context	config>service>vprn>l2tp>group>tunnel
Description	This command configures the local address.
Parameters	<i>ip-address</i> — Specifies the IP address used during L2TP authentication.

local-name

Syntax	local-name <i>host-name</i> no local-name
Context	config>service>vprn>l2tp>group config>service>vprn>l2tp>group>tunnel
Description	This command creates the local host name used by this system for the tunnels in this L2TP group during the authentication phase of tunnel establishment. It can be used to distinguish tunnels. The no form of the command removes the name from the configuration.
Default	no local-name
Parameters	<i>host-name</i> — Specifies the host name, up to 64 characters in length, that the router will use to identify itself during L2TP authentication.

max-retries-estab

Syntax	max-retries-estab <i>max-retries</i> no max-retries-estab
Context	config>service>vprn>l2tp>group config>service>vprn>l2tp>group>tunnel
Description	This command configures the number of retries allowed for this L2TP tunnel while it is established, before its control connection goes down. The no form of the command removes the value from the configuration.

Default	no max-retries-estab
Parameters	<i>max-retries</i> — Specifies the maximum number of retries for an established tunnel.
Values	2 to 7

max-retries-not-estab

Syntax	max-retries-not-estab <i>max-retries</i> no max-retries-not-estab
Context	config>service>vprn>l2tp>group config>service>vprn>l2tp>group>tunnel
Description	This command configures the number of retries allowed for this L2TP tunnel while it is not established, before its control connection goes down. The no form of the command removes the value from the configuration.
Default	no max-retries-not-estab
Parameters	<i>max-retries</i> — Specifies the maximum number of retries for non-established tunnels.
Values	2 to 7

password

Syntax	password <i>password</i> [hash hash2] no password
Context	config>service>vprn>l2tp>group config>service>vprn>l2tp>group>tunnel config>service>vprn>l2tp>group>l2tpv3 config>service>vprn>l2tp>l2tpv3
Description	This command configures the password between L2TP LAC and LNS The no form of the command removes the password.
Default	no password
Parameters	<i>password</i> — Configures the password used for challenge/response calculation and AVP hiding. The maximum length can be up to 20 characters if unhashed, 32 characters if hashed, 54 characters if the hash2 keyword is specified. hash — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

ppp

Syntax	ppp
Context	config>service>vprn>l2tp>group
Description	This command configures PPP for the L2TP tunnel group.

authentication

Syntax	authentication {chap pap pref-chap pref-pap}
Context	config>service>vprn>l2tp>group>ppp
Description	This command configures the PPP authentication protocol to negotiate.

authentication-policy

Syntax	authentication-policy <i>auth-policy-name</i> no authentication-policy
Context	config>service>vprn>l2tp>group>ppp
Description	This command configures the authentication policy.
Parameters	<i>auth-policy-name</i> — Specifies the authentication policy name up to 32 characters in length.

default-group-interface

Syntax	default-group-interface <i>ip-int-name</i> <i>service-id</i> <i>service-id</i> default-group-interface <i>ip-int-name</i> <i>service-name</i> <i>svc-name</i> no default-group-interface
Context	config>service>vprn>l2tp>group>ppp
Description	This command configures the default group interface.
Parameters	<i>ip-int-name</i> — Specifies the interface name up to 32 characters in length.

service-id *service-id* — Specifies the service.

This variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **default-group-interface** *ip-int-name* **service-name** *svc-name* variant can be used in all configuration modes.

Values {*id* | *svc-name*}

id: 1 to 2147483647

svc-name: up to 64 characters (*svc-name* is an alias for input only. The *svc-name* gets replaced with an id automatically by SR OS in the configuration).

service-name *svc-name* — Specifies the service name (instead of service ID) up to 64 characters in length.

keepalive

Syntax	keepalive <i>seconds</i> [hold-up-multiplier <i>multiplier</i>] no keepalive
Context	config>service>vprn>l2tp>group>ppp
Description	This command configures the PPP keepalive interval and multiplier.
Parameters	<i>seconds</i> — Specifies in seconds the interval. Values 10 to 300 <i>multiplier</i> — Specifies the multiplier. Values 1 to 5

lcp-force-ack-accm

Syntax	[no] lcp-force-ack-accm
Context	config>service>vprn>l2tp>group>ppp config>service>vprn>l2tp>group>tunnel>ppp
Description	This command enables or disables the LCP Asynchronous Control Character Map (ACCM) configuration option. When the ACCM configuration option is enabled, the option is acknowledged during the LCP negotiation between the LNS and the PPP client, but no ACCM mapping is performed. By default, the ACCM configuration option is rejected. The no form of this command reverts to the default value.
Default	no lcp-force-ack-accm

mtu

Syntax	mtu <i>mtu-bytes</i> no mtu
Context	config>service>vprn>l2tp>group>ppp
Description	This command configures the maximum PPP MTU size.
Parameters	<i>mtu-bytes</i> — Specifies, in bytes, the maximum PPP MTU size.
Values	512 to 9212

proxy-authentication

Syntax	[no] proxy-authentication
Context	config>service>vprn>l2tp>group>ppp
Description	This command configures the use of the authentication AVPs received from the LAC.

proxy-lcp

Syntax	[no] proxy-lcp
Context	config>service>vprn>l2tp>group>ppp
Description	This command configures the use of the proxy LCP AVPs received from the LAC.

user-db

Syntax	user-db <i>local-user-db-name</i> no user-db
Context	config>service>vprn>l2tp>group>ppp
Description	This command configures the local user database to use for PPP PAP/CHAP authentication.
Parameters	<i>local-user-db-name</i> — Specifies the local user database name.
Values	32 chars max

session-assign-method

Syntax	session-assign-method {existing-first weighted weighted-random} no session-assign-method

Context	config>service>vprn>l2tp>group
Description	This command specifies how new sessions are assigned to one of the set of suitable tunnels that are available or could be made available.
Default	session-assign-method existing-first
Parameters	existing-first — All new sessions are placed by preference in existing tunnels. weighted — Enables weighted preference to tunnels in the group. weighted-random — Enhances the weighted algorithm so that when there are multiple tunnels with an equal number of sessions (equal weight), LAC randomly selects a tunnel.

session-limit

Syntax	session-limit session-limit session-limit unlimited no session-limit
Context	config>service>vprn>l2tp>group config>service>vprn>l2tp>group>tunnel
Description	This command configures the session limit. The value controls how many L2TP session will be allowed within a given context (system, group, tunnel). The no form of the command removes the value from the configuration.
Default	no session-limit
Parameters	<i>session-limit</i> — Specifies the allowed number of sessions within the given context. Values 1 to 131071 unlimited — Specifies the use of the maximum available number of sessions allowed.

3.8.2.3.1 Router L2TP Tunnel Commands

tunnel

Syntax	tunnel tunnel-name [create] no tunnel tunnel-name
Context	config>service>vprn>l2tp>group

Description	This command configures an L2TP tunnel. A tunnel exists between a LAC-LNS pair and consists of a Control Connection and zero or more L2TP sessions. The tunnel carries encapsulated PPP datagrams and control messages between the LAC and the L2TP Network Server (LNS).
Parameters	<i>tunnel-name</i> — Specifies a valid string to identify a L2TP up to 32 characters in length. create — Mandatory while creating a new tunnel.

auto-establish

Syntax	[no] auto-establish
Context	config>service>vprn>l2tp>group>tunnel
Description	This command specifies if this tunnel is to be automatically set up by the system.
Default	no auto-establish

avp-hiding

Syntax	avp-hiding {never sensitive always} no avp-hiding
Context	config>service>vprn>l2tp>group>tunnel
Description	This command configures Attribute Value Pair (AVP) hiding. This capability can be used to avoid the passing of sensitive data, such as user passwords, as clear text in an AVP.



Caution: Nokia recommends that sensitive information not be sent in clear text.

The **no** form of the command removes the parameter of the configuration and indicates that the value on group level will be taken.

Default	no avp-hiding
Parameters	<i>avp-hiding</i> — Specifies the method to be used for the authentication of the tunnel. Values never — AVP hiding is not used. sensitive — AVP hiding is used only for sensitive information (such as username/password). always — AVP hiding is always used.

challenge

Syntax	challenge { always never } no challenge
Context	config>service>vprn>l2tp>group>tunnel
Description	This command configures the use of challenge-response authentication. The no form of the command removes the parameter from the configuration and indicates that the value on group level will be taken.
Default	no challenge
Parameters	always — Specifies that challenge-response authentication should always be used for the tunnel. never — Specifies that challenge-response authentication should never be used for the tunnel.

hello-interval

Syntax	hello-interval <i>hello-interval</i> hello-interval infinite no hello-interval
Context	config>service>vprn>l2tp>group>tunnel
Description	This command configures the number of seconds between sending Hellos for a L2TP tunnel. The no form removes the parameter from the configuration and indicates that the value on group level will be taken.
Parameters	<i>hello-interval</i> — Specifies the time interval, in seconds, between two consecutive tunnel Hello messages. Values 60 to 3600 infinite — Specifies that no Hello messages are sent.

idle-timeout

Syntax	idle-timeout <i>idle-timeout</i> idle-timeout infinite no idle-timeout
Context	config>service>vprn>l2tp>group>tunnel
Description	This command configures the idle timeout to wait before being disconnect.

The **no** form indicates that the parameter will be removed from the configuration and that the value specified on group level will be taken.

Parameters *idle-timeout* — Specifies the idle timeout, in seconds.

Values 0 to 3600

infinite — Specifies that the tunnel will not be closed when idle.

peer

Syntax **peer** *ip-address*
no peer

Context config>service>vprn>l2tp>group>tunnel

Description This command configures the peer address.

The **no** form of the command removes the IP address from the tunnel configuration.

Default no peer

Parameters *ip-address* — Sets the LNS IP address for the tunnel.

preference

Syntax **preference** *preference*
no preference

Context config>service>vprn>l2tp>group>tunnel

Description This command configures a preference number that indicates the relative preference assigned to a tunnel when using a weighted session assignment.

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

Default no preference

Parameters *preference* — Specifies the tunnel preference number with its group. The value 0 corresponds to the highest preference.

Values 0 to 16777215

remote-name

Syntax **remote-name** *host-name*
no remote-name

Context	config>service>vprn>l2tp>group>tunnel
Description	This command configures a string to be compared to the host name used by the tunnel peer during the authentication phase of tunnel establishment.
Parameters	<i>host-name</i> — Specifies a remote host name for the tunnel up to 64 characters in length.

3.8.2.4 Router DHCP Configuration Commands

dhcp

Syntax	dhcp
Context	config>service>vprn
Description	This command enters the context to configure DHCP parameters.

dhcp6

Syntax	dhcp6
Context	config>service>vprn
Description	This command enters the context to configure DHCP6 parameters.

local-dhcp-server

Syntax	local-dhcp-server <i>server-name</i> [create] no local-dhcp-server <i>server-name</i>
Context	config>service>vprn>dhcp config>service>vprn>dhcp6 config>service>vprn>if config>service>vprn>nw-if
Description	This command instantiates a local DHCP server. A local DHCP server can serve multiple interfaces but is limited to the routing context it was which it was created.
Parameters	<i>server-name</i> — Specifies the name of local DHCP server. create — Creates the server name. The create keyword requirement can be enabled/disabled in the environment>create context.

failover

Syntax	failover
Context	config>service>vprn>dhcp
Description	This command enters the context to configure failover parameters.

ignore-mclt-on-takeover

Syntax	ignore-mclt-on-takeover no ignore-mclt-on-takeover
Context	config>service>vprn>dhcp>server>failover config>router>dhcp6>server>failover config>router>dhcp6>server>pool config>service>vprn>dhcp6>server>failover config>service>vprn>dhcp6>server>pool
Description	With this flag enabled, the 'remote' IP address/prefix can be taken over immediately upon entering the PARTNER-DOWN state of the intercommunication link, without having to wait for the MCLT to expire. By setting this flag, the lease times of the existing DHCP clients, while the intercommunication link is in the PARTNER-DOWN state, will still be reduced to the MCLT over time and all new lease times will be set to MCLT. This behavior remain the same as originally intended for MCLT.

Some deployments require that the 'remote' IP address/prefix range starts delegating new IP addresses/prefixes upon the failure of the intercommunication link, without waiting for the intercommunication link to transition from the COMM-INT state into the PARTNER-DOWN state and the MCLT to expire while in PARTNER-DOWN state.

This can be achieved by enabling the ignore-mclt-on-takeover flag and by configuring the partner-down-delay to 0.

Enabling this functionality must be exercised with caution. One needs to keep in mind that the partner-down-delay and MCLT timers were originally introduced to prevent IP address duplication in cases where DHCP redundant nodes transition out-of-sync due to the failure of intercommunication link. These timers (partner-down-delay and MCLT) would ensure that during their duration, the new IP addresses/prefixes are delegated only from one node – the one with local IP address-range/prefix. The drawback is of course that the new IP address delegation is delayed and thus service is impacted.

But if one could ensure that the intercommunication link is always available, then the DHCP nodes would stay in sync and the two timers would not be needed. This is why it is of utmost importance that in this mode of operation, the intercommunication link is well protected by providing multiple paths between the two DHCP nodes. The only event that should cause intercommunication link to fail is the entire nodal failure. This failure is acceptable since in this case only one DHCP node is available to provide new IP addresses/prefixes.

Default no ignore-mclt-on-takeover

maximum-client-lead-time

Syntax **maximum-client-lead-time** [*hrs hours*] [*min minutes*] [*sec seconds*]
no maximum-client-lead-time

Context config>service>vprn>dhcp>server>failover
config>service>vprn>dhcp>server>pool
config>router>vprn>dhcp6>server>failover
config>router>vprn>dhcp6>server>pool
config>service>vprn>dhcp6>server>failover
config>service>vprn>dhcp6>server>pool

Description The maximum-client-lead-time (MCLT) is the maximum time that a DHCP server can extend client's lease time beyond the lease time currently known by the DHCP partner node. In dual-homed environment, the initial lease time for all DHCP clients is by default restricted to MCLT. Consecutive DHCP renewals are allowed to extend the lease time beyond the MCLT.

The MCLT is a safeguard against IP address/prefix duplication in cases of a lease synchronization failure when local-remote failover model is deployed

Once the intercommunication link failure between the redundant DHCP servers is detected, the DHCP IP address range configured as remote will not be allowed to start delegating new leases until the MCLT + partner-down-delay intervals expire. This is to ensure that the new lease that was delegated from the 'local' IP address-range/prefix on one node, but was never synchronized due to the intercommunication link failure, will expire before the same IP address/prefix is allocated from the remote IP address-range/prefix on the other node.

However, the already existing (and synchronized) lease times can be renewed from the remote IP address range at any time, regardless of the state of the intercommunication link (operational or failed).

Lease synchronization failure can be caused either by a node failure, or a failure of the link over which the DHCP leases are synchronized (intercommunication link). Synchronization failure detection can take up to 3 seconds.

During the failure, the DHCP lease time for the new clients will be restricted to MCLT while for the existing clients the lease time will over time (by consecutive DHCP renewals) be gradually reduced to the MCLT.

Default maximum-client-lead-time min 10

Parameters **hrs** *hours* — Specifies the hour parameter of the MCLT.

Values 1 to 23

min *minutes* — Specifies the minute parameter of the MCLT.

Values 1 to 59

sec seconds — Specifies the seconds parameter of the MCLT.

Values 1 to 59

partner-down-delay

Syntax	partner-down-delay [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no partner-down-delay
Context	config>service>vprn>dhcp>server>failover config>router>dhcp6>server>failover config>router>dhcp6>server>pool config>service>vprn>dhcp6>server>failover config>service>vprn>dhcp6>server>pool
Description	Since the DHCP lease synchronization failure can be caused by the failure of the intercommunication link (and not necessary the entire node), there is a possibility the redundant DHCP servers become isolated in the network. In other words, they can serve DHCP clients but they cannot synchronize the lease. This can lead to duplicate assignment of IP addresses, since the servers have configured overlapping IP address ranges but they are not aware of each other's leases. The purpose of the partner-down-delay is to prevent the IP lease duplication during the intercommunication link failure by not allowing new IP addresses to be assigned from the remote IP address range. This timer is intended to provide the operator with enough time to remedy the failed situation and to avoid duplication of IP addresses/prefixes during the failure. During the partner-down-delay time, the prefix designated as remote will be eligible only for renewals of the existing DHCP leases that have been synchronized by the peering node. Only after the sum of the partner-down-delay and the maximum-client-lead-time will the prefix designated as remote be eligible for delegation of the new DHCP leases. When this occurs, we say that the remote IP address range has been taken over. It is possible to expedite the takeover of a remote IP address range so that the new IP leases can start being delegated from that range shortly after the intercommunication failure is detected. This can be achieved by configuring the partner-down-delay timer to 0 seconds, along with enabling the ignore-mclt-on-takeover CLI flag. Caution must be taken before enabling this functionality. It is safe to bypass safety timers (partner-down-delay + MCLT) only in cases where the operator is certain that the intercommunication between the nodes has failed due to the entire node failure and not due to the intercommunication (MCS) link failure. Failed intercommunication due to the nodal failure would ensure that only one node is present in the network for IP address delegation (as opposed to two isolated nodes with overlapping IP address ranges where address duplication can occur). For this reason, the operator must ensure that there are redundant paths between the nodes to ensure uninterrupted synchronization of DHCP leases. In access-driven mode of operation, partner-down-delay has no effect.

Default	partner-down-delay hrs 23 min 59 sec 59
Parameters	hrs <i>hours</i> — Specifies the hour parameter of the partner down delay feature. Values 1 to 23 min <i>minutes</i> — Specifies the minute parameter of the partner down delay feature. Values 1 to 59 sec <i>seconds</i> — Specifies the seconds parameter of the partner down delay feature. Values 1 to 59

peer

Syntax	peer <i>ip-address</i> tag <i>sync-tag-name</i> no peer <i>ip-address</i>
Context	config>service>vprn>dhcp>server>failover config>router>dhcp6>server>failover config>router>dhcp6>server>pool config>service>vprn>dhcp6>server>failover config>service>vprn>dhcp6>server>pool
Description	DHCP leases can be synchronized per DHCP server of DHCP pool. The pair of synchronizing servers or pools is identified by a tag. The synchronization information is carried over the Multi-Chassis Synchronization (MCS) link between the two peers. MCS link is a logical link (IP, or MPLS). MCS runs over TCP, port 45067 and it is using either data traffic or keepalives to detect failure on the communication link between the two nodes. In the absence of any MCS data traffic for more than 0.5sec, MCS will send its own keepalive to the peer. If a reply is not received within 3sec, MCS will declare its operation state as DOWN and the DB Sync state as out-of-sync. MCS will consequently notify its clients (DHCP Server being one of them) of this. It can take up to 3 seconds before the DHCP client realizes that the inter-chassis communication link has failed. The inter-chassis communication link failure does not necessarily assume the same failed fate for the access links. In other words the two redundant nodes can become isolated from each other in the network. This would occur in cases where only the intercommunication (MCS) link fails. It is of utmost importance that this MCS link be highly redundant.
Parameters	<i>ip-address</i> — Specifies the IPv4 address of the peer. sync-tag <i>sync-tag</i> — Specifies a synchronization tag to be used while synchronizing DHCP server or pools.

startup-wait-time

Syntax	[no] startup-wait-time [min <i>minutes</i>] [sec <i>seconds</i>]
Context	config>service>vprn>dhcp6>server>failover config>service>vprn>dhcp6>server>pool config>router>dhcp6>server>failover config>router>dhcp6>server>pool
Description	This command enables startup-wait-time during which each peer waits after the initialization process before assuming the active role for the prefix designated as local or access-driven. This is to avoid transient issues during the initialization process.
Default	startup-wait-time min 2
Parameters	min <i>minutes</i> — Specifies the minute parameter of the startup wait time feature. Values 1 to 10 sec <i>seconds</i> — Specifies the seconds parameter of the startup wait time feature. Values 1 to 59

ignore-rapid-commit

Syntax	[no] ignore-rapid-commit
Context	config>service>vprn>dhcp6>local-dhcp-server
Description	This command specifies whether the Rapid Commit Option (RCO) sent by the DHCPv6 client is processed. If enabled and the client has included an RCO in the solicit, the server ignores the option and processes the remainder of the message as if no RCO were present. The no form of the command disables ignore-rapid-commit.

lease-hold-time

Syntax	lease-hold-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no lease-hold-time
Context	config>service>vprn>dhcp6>server
Description	This command configures the time to remember this lease.
Parameters	[days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] — The lease hold time. Values days: 0 to 3650

hours:	0 to 23
minutes:	0 to 59
seconds:	0 to 5

force-renews

Syntax	[no] force-renews
Context	config>service>vprn>dhcp>server
Description	This command enables the sending of sending force-renew messages. The no form of the command disables the sending of force-renew messages.
Default	no force-renews

pool

Syntax	pool <i>pool-name</i> [create] no pool <i>pool-name</i>
Context	config>service>vprn>dhcp>server
Description	This command configures a DHCP address pool on the router.
Parameters	<i>pool name</i> — Specifies the name of this IP address pool. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. create — Keyword used to create the entity. The create keyword requirement can be enabled/disabled in the environment>create context.

max-lease-time

Syntax	max-lease-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no max-lease-time
Context	config>service>vprn>dhcp>server>pool
Description	This command configures the maximum lease time. The no form of the command returns the value to the default.
Default	max-lease-time days 10

Parameters *time* — Specifies the maximum lease time.

Values

days:	0 to 3650
hours	0 to 23
minutes:	0 to 59
seconds	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>vprn>dhcp>server>pool

Description This command configures the minimum lease time.
 The **no** form of the command returns the value to the default.

Default min-lease-time min 10

Parameters *time* — Specifies the minimum lease time.

Values

days:	0 to 3650
hours	0 to 23
minutes:	0 to 59
seconds	0 to 59

minimum-free

Syntax **minimum-free** *minimum-free* [**percent**] [**event-when-depleted**]
no minimum-free

Context config>service>vprn>dhcp>server>pool

Description This command configures the minimum number of free addresses.
 The **no** form of the command reverts to the default.

Default minimum-free 1

Parameters *minimum-free* — Specifies the desired minimum number of free addresses in this pool. If the actual number of free addresses in this pool falls below this configured minimum, a notification is generated.

Values 0 to 255

percent — indicates the value indicates a percentage.

event-when-depleted — Enables a system-generate event when all available addresses in the pool/subnet of local DHCP server are depleted.

offer-time

Syntax	offer-time [<i>min minutes</i>] [<i>sec seconds</i>] no offer-time
Context	config>service>vprn>dhcp>server>pool
Description	This command configures the offer time. The no form of the command returns the value to the default.
Default	offer-time min 1
Parameters	<i>time</i> — Specifies the offer time. Values minutes: 0 to 10 seconds: 0 to 59

options

Syntax	options
Context	config>service>vprn>dhcp>server>pool
Description	This command enters the context to configure pool options. The options defined here can be overruled by defining the same option in the local user database.

custom-option

Syntax	custom-option <i>option-number</i> address <i>ip-address</i> [<i>ip-address</i>] (DHCP only) custom-option <i>option-number</i> address <i>ipv6-address</i> [<i>ipv6-address</i>] (DHCP6 only) custom-option <i>option-number</i> hex <i>hex-string</i> custom-option <i>option-number</i> string <i>ascii-string</i> no custom-option <i>option-number</i>
Context	config>service>vprn>dhcp>server>pool>options config>service>vprn>dhcp>server>pool>subnet>options
Description	This command configures specific DHCP options. The options defined here can overrule options in the local user database. The no form of the removes the option from the configuration.

Parameters	<i>option-number</i> — Specifies the option number that the DHCP server uses to send the identification strings to the DHCP client. Values 1 to 254 address <i>ip-address</i> — Specifies the IP address of this host. Up to 4 IP addresses can be specified in a single statement. <i>hex hex-string</i> — Specifies the hex value of this option. Values 0x0 to 0xFFFFFFFF...(maximum 254 hex nibbles) string <i>ascii-string</i> — Specifies the value of this option. Values up to 127 characters maximum

dns-server

Syntax	dns-server <i>ip-address</i> [<i>ip-address</i>] (DHCP only) dns-server <i>ipv6-address</i> [<i>ipv6-address</i>] (DHCP6 only) no dns-server
Context	config>service>vprn>dhcp>server>pool>options
Description	This command configures the IP address of the DNS server.
Parameters	<i>ip-address</i> — The IP address of the DNS server. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 to 223.255.255.255 (with support of /31 subnets). Up to 4 IP addresses maximum.

domain-name

Syntax	domain-name <i>domain-name</i> no domain-name
Context	config>service>vprn>dhcp>server>pool>options
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 the command removes the name from the configuration.
Parameters	<i>domain-name</i> — Specifies the domain name for the client. Values Up to 127 characters

renew-timer

Syntax	renew-timer [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no renew-timer		
Context	config>service>vprn>dhcp6>server>pool>prefix		
Description	This command configures the renew-timer (T1), the time at which the client contacts the server from which the addresses in the IA_NA or IA_PD were obtained to extend the lifetimes of the addresses or prefixes assigned to the client.		
Default	renew-timer min 30		
Parameters	<i>time</i> — Specifies the renew timer.		
	Values		
		days:	0 to 7
		hours	0 to 23
		minutes:	0 to 59
		seconds	0 to 59

rebind-timer

Syntax	rebind-timer [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no rebind-timer		
Context	config>service>vprn>dhcp6>server>pool>prefix		
Description	This command configures the rebind-timer (T2), the time at which the client contacts any available server to extend the lifetimes of the addresses or prefixes assigned to the client.		
Default	rebind-timer min 48		
Parameters	[days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] — Specifies the rebind timer.		
	Values		
		days	0 to 14
		hours	0 to 23
		minutes	0 to 59
		seconds	0 to 59

prefix

Syntax	prefix <i>ipv6-address/prefix-length</i> [failover {local remote access-driven}] [pd] [wan-host] [create] no prefix <i>ipv6-address/prefix-length</i>

Context	config>router>dhcp6>server>pool config>service>vprn>dhcp6>server>pool
Description	This command allows a list of prefixes (using the prefix command multiple times) to be routed to hosts associated with this pool. Each prefix will be represented in the associated FIB with a reference to the pool. Prefixes are defined as being for prefix delegation (pd) or use on a WAN interface or host (wan-host).
Default	failover local
Parameters	<i>ipv6-address</i> — Specifies the 128-bit IPv6 address. Values 128-bit hexadecimal IPv6 address in compressed form <i>prefix-length</i> — Specifies the length of any associated aggregate prefix. Values 32 to 63 failover — This command designates a IPv6 prefix as local, remote or access-driven. This is used when multi-chassis synchronization is enabled. local — An IPv6 prefix designated as local is used for new lease grants or to renew the existing lease grants. Local prefix designation should be always paired with the remote designation of the same prefix on the peering node. The IPv6 prefix configured as local on one node can only be configured as remote on the other node. No other combination is allowed between the two nodes for an IPv6 prefix that is configured as local. The dhcpv6 relay could point to both IPv6 DHCP server addresses— the one hosting the local IPv6 prefix and the one hosting the corresponding remote IPv6 prefix. Under normal circumstances the new lease will always be allocated from the local IPv6 prefix while the leases can be renewed from either IPv6 prefix (local or remote). Under network failure, the remote IPv6 prefix can be taken over according to the intercommunication link state transitions and associated timers. remote — An IPv6 prefix designated as remote is used only to renew the existing DHCP leases. The new leases will be delegated from it only after the maximum-client-lead-time + partner-down-delay time elapses. At that point we say that the remote IPv6 prefix has been taken over. To ensure faster takeover, the partner-down-delay can be set to 0 and the MCLT time can be ignored. Extra caution should be exercised when enabling this mode of operation, as described in the configuration guides. The IPv6 prefix configured as remote on one node can only be configured as local on the other node. No other combination is allowed between the two nodes for an IP address ranges that is configured as remote. access-driven — An IPv4 prefix designated as access-driven is used for new lease grants or to renew the existing lease grants regardless of the state of the intercommunication link (operational or failed). In this mode of operation the IPv6 prefix is actively shared between the two DHCPv6 server nodes. This can be used on both DHCPv6 servers only in cases where the access protection mechanism (SRRP or MC-LAG) will ensure that there is only a single active path for DHCPv6 clients using the same IPv6 prefix available to one of the redundant DHCPv6 nodes.

The IPv6 prefix configured as access-driven on one node can only be configured as access-driven on the other node. No other combination is allowed between the two nodes for an IPv6 prefix that is configured as access-driven.

There must be no crosslinks between the DHCPv6 servers that have IPv6 address ranges configured in access-driven failover mode. In other words, each node must have the dhcp-relay pointing to the IPv6 address of the local DHCPv6 server. This IPv6 address must be the same on both nodes. For example, both DHCPv6 servers should have a loopback address configured with the same IPv6 address (IPv4 or IPv6) and a DHCPv6 server associated with this loopback address. Those IPv6 addresses must not be advertised outside of each box. The DHCPv6 relay in each node would point to its local DHCPv6 server via this loopback IPv6 address.

pd — Specifies that this aggregate is used by IPv6 ESM hosts for DHCPv6 prefix-delegation.

wan-host — Specifies that this aggregate is used by IPv6 ESM hosts for local addressing or by a routing gateway's WAN interface.

preferred-lifetime

Syntax	preferred-lifetime [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no preferred-lifetime		
Context	config>service>vprn>dhcp6>server>pool>prefix		
Description	The preferred lifetime for the IPv6 prefix or address in the option, expressed in units of seconds. When the preferred lifetime expires, any derived addresses are deprecated.		
Default	preferred-lifetime hrs 1		
Parameters	<i>time</i> — Specifies the preferred lifetime.		
	Values		
	days	0 to 3650	
	hours	0 to 23	
	minutes	0 to 59	
	seconds	0 to 59	

valid-lifetime

Syntax	valid-lifetime [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no valid-lifetime		
Context	config>service>vprn>dhcp6>server>pool>prefix		
Description	The valid lifetime for the IPv6 prefix or address in the option, expressed in units of seconds.		
Default	valid-lifetime days 1		

Parameters	<i>time</i> — Specifies the valid lifetime.		
	Values		
		days	0 to 3650
		hours	0 to 23
		minutes	0 to 59
		seconds	0 to 59

use-link-address

Syntax	use-link-address [<i>scope scope</i>] no use-link-address
Context	config>service>vprn>dhcp6>server
Description	This command specifies whether the GI address selects a single subnet or a pool. The no form of the command reverts to the default.
Default	subnet
Parameters	scope <i>scope</i> — Specifies the scope of the IP address selection.
	Values subnet, pool

user-ident

Syntax	user-ident <i>user-ident</i> no user-ident
Context	config>service>vprn>dhcp6>server
Description	This command specifies which method is used by the local DHCP server to uniquely identify a user. The no form of the command reverts to the default.
Default	user-ident duid
Parameters	<i>user-ident</i> — Configures the user identification method.
	Values duid, interface-id, interface-id-link-local

lease-rebind-time

Syntax	lease-rebind-time [<i>days days</i>] [<i>hrs hours</i>] [<i>min minutes</i>] [<i>sec seconds</i>] no lease-rebind-time

Context	config>service>vprn>dhcp>server>pool>options		
Description	This command configures the time the client transitions to a rebinding state. The no form of the command removes the time from the configuration.		
Parameters	<i>time</i> — Specifies the lease rebind time.		
	Values		
		days	0 to 3650
		hours	0 to 23
		minutes	0 to 59
		seconds	0 to 59

lease-renew-time

Syntax	lease-renew-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no lease-renew-time		
Context	config>service>vprn>dhcp>server>pool>options		
Description	This command configures the time the client transitions to a renew state. The no form of the command removes the time from the configuration.		
Parameters	<i>time</i> — Specifies the lease renew time.		
	Values		
		days	0 to 3650
		hours	0 to 23
		minutes	0 to 59
		seconds	0 to 59

lease-time

Syntax	lease-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] no lease-time		
Context	config>service>vprn>dhcp>server>pool>options		
Description	This command configures the amount of time that the DHCP server grants to the DHCP client permission to use a particular IP address. The no form of the command removes the lease time parameters from the configuration.		

Parameters *time* — Specifies the lease time.

Values

days	0 to 3650
hours	0 to 23
minutes	0 to 59
seconds	0 to 59

netbios-name-server

Syntax **netbios-name-server** *ip-address* [*ip-address...*(up to 4 max)]
no netbios-name-server

Context config>service>vprn>dhcp>server>pool>options

Description This command configures Network Basic Input/Output System (NetBIOS) name server IP addresses.

Parameters *ip-address* — The IP address of the NetBIOS name server. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 to 223.255.255.255 (with support of /31 subnets). Up to 4 IP addresses maximum.

netbios-node-type

Syntax **netbios-node-type** *netbios-node-type*
no netbios-node-type

Context config>service>vprn>dhcp>server>pool>options

Description This command configures the Network Basic Input/Output System (NetBIOS) node type.

Parameters *netbios-node-type* — Specifies the netbios node type.

Values

- B — Broadcast node uses broadcasting to query nodes on the network for the owner of a NetBIOS name.
- P — Peer-to-peer node uses directed calls to communicate with a known NetBIOS name server for the IP address of a NetBIOS machine name.
- M — Mixed node uses broadcast queries to find a node, and if that fails, queries a known P-node name server for the address.
- H — Hybrid node is the opposite of the M-node action so that a directed query is executed first, and if that fails, a broadcast is attempted.

server

Syntax	server <i>server-name</i> no server
Context	config>service>ies>sub-if>grp-if>local-address-assignment config>service>ies>sub-if>local-address-assignment config>service>vprn>sub-if>grp-if>local-address-assignment config>service>vprn>sub-if>local-address-assignment
Description	This command designates a local DHCPv4 server for local pools management where IPv4 addresses for PPPoXv4 clients will be allocated without the need for the internal DHCP relay-agent. Those addresses will be tied to PPPoX sessions and they will be de-allocated when the PPPoX session is terminated.
Parameters	<i>server-name</i> — Specifies the name of the local DHCP server.

client-application

Syntax	client-application [ppp-v4] no client-application
Context	config>service>ies>sub-if>grp-if>local-address-assignment config>service>ies>sub-if>local-address-assignment config>service>vprn>sub-if>grp-if>local-address-assignment config>service>vprn>sub-if>local-address-assignment
Description	This command enables local DHCP Server pool management for PPPoXv4 clients. A pool of IP addresses can be shared between IPE clients that rely on DHCP protocol (lease renewal process) and PPPoX clients where address allocation is not dependent on DHCP messaging but instead an IP address allocation within the pool is tied to the PPPoX session.

default-pool

Syntax	default-pool <i>pool-name</i> no default-pool
Context	config>service>ies>sub-if>grp-if>local-address-assignment config>service>ies>sub-if>local-address-assignment config>service>vprn>sub-if>grp-if>local-address-assignment config>service>vprn>sub-if>local-address-assignment
Description	This command references a default DHCP address pool for local PPPoX pool management in case that the pool-name is not returned via Radius or LUDB.
Parameters	<i>pool-name</i> — Specifies the name of the local DHCP server pool.

subnet

Syntax	subnet { <i>ip-address/mask</i> <i>ip-address netmask</i> } [create] no subnet { <i>ip-address/mask</i> <i>ip-address netmask</i> }
Context	config>service>vprn>dhcp>server>pool
Description	This command creates a subnet of IP addresses to be served from the pool. The subnet cannot include any addresses that were assigned to subscribers without those addresses specifically excluded. When the subnet is created no IP addresses are made available until a range is defined.
Parameters	<i>ip-address</i> — Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 to 223.255.255.255 (with support of /31 subnets). <i>mask</i> — The subnet mask in dotted decimal notation. Allowed values are dotted decimal addresses in the range 128.0.0.0 to 255.255.255.252. A mask of 255.255.255.255 is reserved for system IP addresses. <i>netmask</i> — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains. create — Keyword used to create the entity. The create keyword requirement can be enabled/disabled in the environment>create context.

address-range

Syntax	address-range <i>start-ip-address end-ip-address</i> [failover { local remote access-driven }] no address-range <i>start-ip-address end-ip-address</i>
Context	config>service>vprn>dhcp>server>pool>subnet config>router>dhcp>server>pool>subnet
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 will be included (other than specific excluded addresses). The only two valid failover combinations between the two redundant DHCP nodes are: • local - remote • access-driven - access-driven
Default	failover local
Parameters	<i>start-ip-address</i> — Specifies the start address of this range to include. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

end-ip-address — Specifies the end address of this range to include. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

failover — This command designates an address range as local, remote or access-driven. This is used when multi-chassis synchronization is enabled.

local — An IPv4 address-range designated as local is used for new lease grants or to renew the existing lease grants. Local address-range designation should be always paired with the remote designation of the same address-range on the peering node. The IP address range configured as local on one node can only be configured as remote on the other node. No other combination is allowed between the two nodes for an IP address ranges that is configured as local.

The dhcp relay could point to both IP DHCP server addresses—the one hosting the local IP address range and the one hosting the corresponding remote IP address range. Under normal circumstances the new lease will always be allocated from the local IP address range while the leases can be renewed from either IP address range (local or remote). Under network failure, the remote IP address range can be taken over according to the intercommunication link state transitions and associated timers.

remote — An IPv4 address-range designated as remote is used only to renew the existing DHCP leases. The new leases will be delegated from it only after the maximum-client-lead-time + partner-down-delay time elapses. At that point we say that the remote IP address range has been taken over.

To ensure faster takeover, the partner-down-delay can be set to 0 and the MCLT time can be ignored. Extra caution should be exercised when enabling this mode of operation, as described in the configuration guides.

The IP address range configured as remote on one node can only be configured as local on the other node. No other combination is allowed between the two nodes for an IP address ranges that is configured as remote.

access-driven — An IPv4 address-range designated as access-driven is used for new lease grants or to renew the existing lease grants regardless of the state of the intercommunication link (operational or failed). In this mode of operation the IP address-range is actively shared between the two DHCP server nodes. This can be used on both DHCP servers only in cases where the access protection mechanism (SRRP or MC-LAG) will ensure that there is only a single active path for DHCP clients using the same IP address range available to one of the redundant DHCP nodes.

The IP address range configured as access-driven on one node can only be configured as access-driven on the other node. No other combination is allowed between the two nodes for an IP address ranges that is configured as access-driven.

There must be no crosslinks between the DHCP servers that have IP address ranges configured in access-driven failover mode. In other words, each node must have the dhcp-relay pointing to the IP address of the local DHCP server. This IP address must be the same on both nodes. For example, both DHCP servers should have a loopback address configured with the same IP address (IPv4 or IPv6) and a DHCP server associated with this loopback address. Those IP addresses must not be advertised outside of each box. The DHCP relay in each node would point to its local DHCP server via this loopback IP address.

drain

Syntax	[no] drain
Context	config>service>vprn>dhcp>server>pool>subnet
Description	This command subnet draining which means no new leases can be assigned from this subnet and existing leases are cleaned up upon renew/rebind. The no form of the command means the subnet is active and new leases can be assigned from it.

exclude-addresses

Syntax	[no] exclude-addresses <i>start-ip-address</i> [<i>end-ip-address</i>]
Context	config>service>vprn>dhcp>server>pool>subnet
Description	This command specifies a range of IP addresses that excluded from the pool of IP addresses in this subnet.
Parameters	<i>start-ip-address</i> — Specifies the start address of this range to exclude. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets). <i>end-ip-address</i> — Specifies the end address of this range to exclude. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

maximum-declined

Syntax	maximum-declined <i>maximum-declined</i> no maximum-declined
Context	config>service>vprn>dhcp>server>pool>subnet

Description	This command configures the maximum number of declined addresses allowed.
Default	maximum-declined 64
Parameters	<i>maximum-declined</i> — Specifies the maximum number of declined addresses allowed.
Values	0 to 4294967295

minimum-free

Syntax	minimum-free <i>minimum-free</i> [percent] [event-when-depleted] no minimum-free
Context	config>service>vprn>dhcp>server>pool>subnet
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	minimum-free 1
Parameters	<i>minimum-free</i> — Specifies the minimum number of free addresses in this subnet.
Values	0 to 255
	percent — indicates the value indicates a percentage
	event-when-depleted — Enables a system-generate event when all available addresses in the pool/subnet of local DHCP server are depleted.

default-router

Syntax	default-router <i>ip-address</i> [<i>ip-address</i>] no default-router
Context	config>service>vprn>dhcp>server>pool>subnet
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 the command removes the address(es) from the configuration.
Parameters	<i>ip-address</i> — Specifies the IP address of the default router. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets). Up to 4 IP addresses maximum.

subnet-mask

Syntax	subnet-mask <i>ip-address</i> no subnet-mask
Context	config>service>vprn>dhcp>server>pool>subnet
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 the command removes the address from the configuration.
Parameters	<i>ip-address</i> — Specifies the IP address of the subnet mask. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 to 223.255.255.255 (with support of /31 subnets).

use-gi-address

Syntax	[no] use-gi-address
Context	config>service>vprn>dhcp>server
Description	This command enables the use of gi-address matching. If the gi-address flag is enabled, a pool can be used even if a subnets is not found. If the local-user-db-name is not used, the gi-address flag is used and addresses are handed out by GI only. If a user must be blocked from getting an address the server maps to a local user database and configures the user with no address. A pool can include multiple subnets. Since the GI is shared by multiple subnets in a subscriber-interface the pool may provide IP addresses from any of the subnets included when the GI is matched to any of its subnets. This allows a pool to be created that represents a sub-int.
Default	no use-gi-address

use-pool-from-client

Syntax	[no] use-pool-from-client
Context	config>service>vprn>dhcp>server config>service>vprn>dhcp6>server
Description	This command specifies if the IP address pool to be used by this server is the pool indicated by the vendor-specific sub-option 13 of the DHCP Option 82. When enabled, the pool indicated by the sub-option 13 is to be used.

The **no** form of the command indicates that the pool selection is specified by the value of **use-gi-address** setting.

user-db

Syntax	user-db <i>local-user-db-name</i> no user-db
Context	config>service>vprn>dhcp>server
Description	This command configures a local user database for authentication.
Default	no user-db
Parameters	<i>local-user-db-name</i> — Specifies the name of a local user database.

3.8.2.5 IGMP Commands

igmp

Syntax	[no] igmp
Context	config>service>vprn
Description	This command enters the context to configure IGMP parameters. The no form of the command disables IGMP.
Default	disabled

group-interface

Syntax	[no] group-interface <i>ip-int-name</i> [no] group-interface fwd-service <i>service-id ip-int-name</i>
Context	config>service>vprn>igmp
Description	This command configures IGMP group interfaces. The no form of the command reverts to the default.
Parameters	<i>ip-int-name</i> — Specifies the name of the IP interface. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

fwd-service *service-id* — Specifies the service ID. This is only configured in the retailer VRF. This construct references the wholesaler service under which the group-interface (and the subscriber) is actually defined.

Values 1 to 2147483650, svc-name up to 64 char maximum

disable-router-alert-check

Syntax [no] **disable-router-alert-check**

Context config>service>ies>sub-if>grp-if>sap>igmp-host-tracking
config>service>vprn>igmp>grp-if
config>service>vprn>igmp>if

Description This command disables the IGMP router alert check option.
The **no** form of the command enables the router alert check.

import

Syntax **import** *policy-name*
no import

Context config>service>vprn>igmp>grp-if
config>service>vprn>igmp>if

Description This command specifies the policy that is to be applied on this interface.

Parameters *policy-name* — Specifies the policy to filter IGMP packets.

max-groups

Syntax **max-groups** *value*
no max-groups

Context config>service>vprn>igmp>grp-if
config>service>vprn>igmp>if

Description This command configures the maximum number of groups for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed.
The **no** form of the command removes the value.

Parameters *value* — Specifies the maximum number of groups for this interface.
Values 1 to 16000

max-sources

Syntax **max-sources** *max-sources*
no max-sources

Context config>service>vprn>igmp>grp-if
 config>service>vprn>igmp>if

Description This command specifies the maximum number of sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of sources, the sources that are already accepted are not deleted. Only new sources will not be allowed.

Parameters *sources* — Specifies the maximum number of sources for this interface.
Values 1 to 1000

max-grp-sources

Syntax **max-grp-sources** *max-group-sources*
no max-grp-sources

Context config>service>vprn>igmp>grp-if
 config>service>vprn>igmp>if

Description This command configures the maximum number of group sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources will not be allowed.

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

Default 0

Parameters **1 to 32000** — Specifies the maximum number of group source.
Values 1 to 32000

mcac

Syntax **mcac**

Context config>service>vprn>igmp>group-interface
config>service>vprn>igmp>interface
config>service>vprn>mld>group-interface
config>service>vprn>mld>interface
config>service>vprn>pim>interface

Description This command configures multicast CAC policy and constraints for this interface.

if-policy

Syntax **if-policy** *mcac-if-policy-name*
no if-policy

Context config>service>vprn>igmp>grp-if>mcac
config>service>vprn>igmp>interface>mcac
config>service>vprn>mld>grp-if>mcac
config>service>vprn>mld>interface>mcac
config>service>vprn>pim>interface>mcac

Description This command assigns existing MCAC interface policy to this interface.

The **no** form of the command removes the MCAC interface policy association.

Default no if-policy

Parameters *mcac-if-policy-name* — Specifies an existing MCAC interface policy.

mc-constraints

Syntax **mc-constraints**

Context config>service>vprn>igmp>grp-if>mcac
config>service>vprn>igmp>interface>mcac
config>service>vprn>mld>grp-if>mcac
config>service>vprn>mld>interface>mcac
config>service>vprn>pim>interface>mcac

Description This command enters the context to configure multicast CAC constraints.

level

Syntax **level** *level-id* **bw** *bandwidth*
no level *level-id*

Context config>service>vprn>igmp>interface>mcac>mc-constraints
config>service>vprn>mld>interface>mcac>mc-constraints

config>service>vprn>pim>interface>mcac>mc-constraints

Description This command configures interface levels and associated bandwidth for multicast CAC policy.

Parameters *level-id* — Specifies an entry for the multicast CAC policy constraint level configured on this system.

Values 1 to 8

bandwidth — Specifies the bandwidth in kb/s for the level.

Values 1 to 2147483647

number-down

Syntax **number-down** *number-lag-port-down* **level** *level-id*
no number-down

Context config>service>vprn>igmp>interface>mcac>mc-constraints
config>service>vprn>mld>interface>mcac>mc-constraints
config>service>vprn>pim>interface>mcac>mc-constraints

Description This command configures the number of ports down and level for interface's multicast CAC policy.

Default not enabled

Parameters *number-lag-port-down* — If the number of ports available in the LAG is reduced by the number of ports configured in this command here then bandwidth allowed for bundle and/or interface will be as per the levels configured in this context.

Values 1 to 64 (for 64-link LAG)
1 to 32 (for other LAGs)

level-id — Specifies an entry for the multicast CAC policy constraint level configured on this system.

Values 1 to 8

use-lag-port-weight

Syntax [**no**] **use-lag-port-weight**

Context config>service>vprn>igmp>interface>mcac>mc-constraints
config>service>vprn>mld>interface>mcac>mc-constraints
config>service>vprn>pim>interface>mcac>mc-constraints

Description This command enables port weight to be used when determining available bandwidth per level when LAG ports go down/come up. The command is required for proper operation on mixed port-speed LAGs and can be used for non-mixed port-speed LAGs as well.

Default	no use-lag-port-weight - port number is used when determining available bandwidth per level when LAG ports go down/come up.

policy

Syntax	policy <i>policy-name</i> no policy
Context	config>service>vprn>igmp>grp-if>mcac config>service>vprn>igmp>interface>mcac config>service>vprn>mld>grp-if>mcac config>service>vprn>mld>interface>mcac config>service>vprn>pim>interface>mcac
Description	This command references the global channel bandwidth definition policy that is used for HMCAC and HQoS Adjust. HQoS Adjustment is supported with redirection enabled or per-host-replication disabled. In other words, the policy from the redirected interface is used for HQoS Adjustment. Hierarchical MCAC (HMCAC) is supported with redirection enabled or per-host-replication disabled. In HMCAC, the subscriber is checked first against its bandwidth limits followed by the check on the redirected interface against the bandwidth limits defined under the redirected interface. In the HMCAC case, the channel definition policy must be referenced under the redirected interface level.
Default	No policy is referenced.
Parameters	<i>policy-name</i> — Specifies the name of the global MCAC channel definition policy defined under the hierarchy config>router>mcac>policy .

policy

Syntax	policy <i>policy-name</i> no policy
Context	config>service>vprn>igmp>if>mcac config>service>vprn>pim>if>mcac
Description	This command configures the multicast CAC policy name.
Parameters	<i>policy-name</i> — The multicast CAC policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

unconstrained-bw

Syntax	unconstrained-bw <i>bandwidth</i> mandatory-bw <i>mandatory-bw</i> no unconstrained-bw
Context	config>service>vprn>igmp>grp-if>mcac config>service>vprn>igmp>interface>mcac config>service>vprn>mld>grp-if>mcac config>service>vprn>mld>interface>mcac config>service>vprn>pim>interface>mcac
Description	This command configures the bandwidth for the interface's multicast CAC policy traffic. When disabled (no unconstrained-bw) there will be no checking of bandwidth constraints on the interface level. When enabled and a policy is defined, enforcement is performed. The allocated bandwidth for optional channels should not exceed the unconstrained-bw minus the mandatory-bw and the mandatory channels have to stay below the specified value for the mandatory-bw . After this interface check, the bundle checks are performed.
Parameters	<i>bandwidth</i> — The bandwidth assigned for the interface's MCAC policy traffic in kb/s. Values 0 to 2147483647 mandatory-bw <i>mandatory-bw</i> — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface, in kb/s. If the <i>bandwidth</i> value is 0, no mandatory channels are allowed. If <i>bandwidth</i> is not configured, then all mandatory and optional channels are allowed. If the value of <i>mandatory-bw</i> is equal to the value of <i>bandwidth</i> , then all the unconstrained bandwidth on a given interface is allocated to mandatory channels configured through multicast CAC policy on that interface and no optional groups (channels) are allowed. The value of <i>mandatory-bw</i> should always be less than or equal to that of <i>bandwidth</i> , An attempt to set the value of <i>mandatory-bw</i> greater than that of <i>bandwidth</i> , will result in inconsistent value error. Values 0 to 2147483647

query-src-ip

Syntax	query-src-ip <i>ip-address</i> no query-src-ip
Context	config>service>vprn>igmp>grp-if
Description	This command configures the query source IP address for the group interface. This IP address overrides the source IP address configured at the router level. The no form of the command removes the IP address.
Parameters	<i>ip-address</i> — Sets the source IPv4 address for all subscriber's IGMP queries.

sub-hosts-only

Syntax	[no] sub-hosts-only
Context	config>service>vprn>igmp>grp-if
Description	This command enables the IGMP traffic from known hosts only. The no form of the command disable the IGMP traffic from known hosts only

subnet-check

Syntax	[no] subnet-check
Context	config>service>vprn>igmp>grp-if
Description	This command enables local subnet checking for IGMP. The no form of the command disables local subnet checking for IGMP.

version

Syntax	version <i>version</i> no version
Context	config>service>vprn>igmp>grp-if
Description	This command configures the version of IGMP. The no form of the command removes the version.
Parameters	<i>version</i> — Specifies the IGMP version. Values 1, 2, or 3

grp-if-query-src-ip

Syntax	grp-if-query-src-ip <i>ip-address</i> no grp-if-query-src-ip
Context	config>service>vprn>igmp
Description	This command configures the query source IP address for all group interfaces. The no form of the command removes the IP address.

interface

Syntax	interface <i>ip-int-name</i> no interface
Context	config>service>vprn>igmp
Description	This command enters the context to configure IGMP interface parameters.
Parameters	<i>ip-int-name</i> — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.
Values	1 to 32 characters maximum

import

Syntax	import <i>policy-name</i> no import
Context	config>service>vprn>igmp>if
Description	This command imports a policy to filter IGMP packets. The no form of the command removes the policy association from the IGMP instance.
Default	no import — No import policy specified.
Parameters	<i>policy-name</i> — The import route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. The specified name(s) must already be defined.

max-groups

Syntax	max-groups <i>value</i> no max-groups
Context	config>service>vprn>igmp>if
Description	This command specifies the maximum number of groups for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed.
Default	0, no limit to the number of groups

Parameters	<i>value</i> — Specifies the maximum number of groups for this interface.
Values	1 to 16000

static

Syntax	static
Context	config>service>vprn>igmp>if
Description	This command tests forwarding on an interface without a receiver host. When enabled, data is forwarded to an interface without receiving membership reports from host members.

group

Syntax	[no] group <i>grp-ip-address</i> [no] group start <i>grp-ip-address</i> end <i>grp-ip-address</i> [step <i>ip-address</i>]
Context	config>service>vprn>igmp>if>static
Description	This command adds a static multicast group either as a (*,G) or one or more (S,G) records. Use IGMP static group memberships to test multicast forwarding without a receiver host. When IGMP static groups are enabled, data is forwarded to an interface without receiving membership reports from host members. When static IGMP group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static IGMP group entries do not generate join messages toward the RP.
Parameters	<i>grp-ip-address</i> — Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group. The address must be in dotted decimal notation. start <i>grp-ip-address</i> — Specifies the start multicast group address. end <i>grp-ip-address</i> — Specifies the end multicast group address. step <i>ip-address</i> — Specifies the step increment.

source

Syntax	source <i>ip-address</i>
Context	config>service>vprn>igmp>if>static>group
Description	This command specifies an IPv4 unicast address that sends data on an interface. This enables a multicast receiver host to signal a router the group is to receive multicast traffic from, and from the sources that the traffic is expected.

The **source** command is mutually exclusive with the specification of individual sources for the same group.

The **source** command in combination with the group is used to create a specific (S,G) static group entry.

Use the **no** form of the command to remove the source from the configuration.

Parameters *ip-address* — Specifies the IPv4 unicast address.

starg

Syntax **starg**

Context config>service>vprn>igmp>if>static>group

Description This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified.

Use the **no** form of the command to remove the starg entry from the configuration.

subnet-check

Syntax [**no**] **subnet-check**

Context config>service>vprn>igmp>if

Description This command enables subnet checking for IGMP messages received on this interface. All IGMP packets with a source address that is not in the local subnet are dropped.

Default enabled

version

Syntax **version** *version*
no version

Context config>service>vprn>igmp>if

Description This command specifies the IGMP version. If routers run different versions of IGMP, they will negotiate the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version. For IGMP to function correctly, all routers on a LAN should be configured to run the same version of IGMP on that LAN.

For IGMPv3, a multicast router that is also a group member performs both parts of IGMPv3, receiving and responding to its own IGMP message transmissions as well as those of its neighbors.

Default	version 3
Parameters	<i>version</i> — Specifies the IGMP version number.
Values	1, 2, 3

query-interval

Syntax	query-interval <i>seconds</i> no query-interval
Context	config>service>vprn>igmp
Description	This command specifies the frequency that the querier router transmits general host-query messages. The host-query messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.
Default	query-interval 125
Parameters	<i>seconds</i> — The time frequency, in seconds, that the router transmits general host-query messages.
Values	2 to 1024

query-last-member-interval

Syntax	query-last-member-interval <i>seconds</i>
Context	config>service>vprn>igmp
Description	This command configures the frequency at which the querier sends group-specific query messages including messages sent in response to leave-group messages. The lower the interval, the faster the detection of the loss of the last member of a group.
Default	query-last-member-interval 1
Parameters	<i>seconds</i> — Specifies the frequency, in seconds, at which query messages are sent.
Values	1 to 1024

query-response-interval

Syntax	query-response-interval <i>seconds</i>
Context	config>service>vprn>igmp
Description	This command specifies how long the querier router waits to receive a response to a host-query message from a host.

Default	query-response-interval 10
Parameters	<i>seconds</i> — Specifies the length of time to wait to receive a response to the host-query message from the host.
Values	1 to 1023

robust-count

Syntax	robust-count <i>robust-count</i> no robust-count
Context	config>service>vprn>igmp
Description	This command configures the robust count. The robust-count variable allows tuning for the expected packet loss on a subnet. If a subnet anticipates losses, the robust-count variable can be increased.
Default	robust-count 2
Parameters	<i>robust-count</i> — Specifies the robust count value.
Values	2 to 10

ssm-translate

Syntax	ssm-translate
Context	config>service>vprn>igmp config>service>vprn>igmp>if
Description	This command enters the context to configure group ranges which are translated to SSM (S,G) entries. If the static entry needs to be created, it has to be translated from a IGMPv1 IGMPv2 request to a Source Specific Multicast (SSM) join. An SSM translate source can only be added if the starg command is not enabled. An error message is generated if you try to configure the source command with starg command enabled.

grp-range

Syntax	[no] grp-range <i>start end</i>
Context	config>service>vprn>igmp>ssm-translate
Description	This command is used to configure group ranges which are translated to SSM (S,G) entries.
Parameters	<i>start</i> — An IP address that specifies the start of the group range.

end — An IP address that specifies the end of the group range. This value should always be greater than or equal to the value of the *start* value.

source

Syntax	[no] source <i>ip-address</i>
Context	config>service>vprn>igmp>ssm-translate>grp-range
Description	This command specifies the source IP address for the group range. Whenever a (*,G) report is received in the range specified by grp-range <i>start</i> and <i>end</i> parameters, it is translated to an (S,G) report with the value of this object as the source address.
Parameters	<i>ip-address</i> — Specifies the IP address that will be sending data.

igmp-host-tracking

Syntax	igmp-host-tracking
Context	config>service>vprn config>service>vprn>sap
Description	This command enters the context to configure IGMP host tracking parameters.

expiry-time

Syntax	expiry-time <i>expiry-time</i> no expiry-time
Context	config>service>vprn>igmp-trk config>service>vprn>sap>igmp-trk
Description	This command configures the time that the system continues to track inactive hosts. The no form of the command removes the values from the configuration.
Default	no expiry-time
Parameters	<i>expiry-time</i> — Specifies the time, in seconds, that this system continues to track an inactive host. Values 1 to 65535

import

Syntax	import <i>policy-name</i>

no import

Context	config>service>vprn>sap>igmp-trk
Description	This command associates an import policy to filter IGMP packets. The no form of the command removes the values from the configuration.
Default	no import
Parameters	<i>policy-name</i> — Specifies the import policy name.

max-num-groups

Syntax	max-num-groups <i>max-num-groups</i> no max-num-groups
Context	config>service>vprn>sub-if>grp-if>sap>igmp-trk
Description	This command configures the maximum number of multicast groups allowed to be tracked. The no form of the command removes the values from the configuration.
Default	no max-num-groups
Parameters	<i>max-num-groups</i> — Specifies the maximum number of multicast groups allowed to be tracked. Values 1 to 196607

max-num-grp-sources

Syntax	max-num-grp-sources [<i>number</i>] no max-num-grp-sources
Context	config>service>vprn>sub-if>grp-if>sap>igmp-host-tracking
Description	This command configures the maximum number of multicast sources allowed to be tracked per group. The no form of the command disables the check.
Default	no max-num-grp-sources
Parameters	<i>number</i> — Specifies the maximum number of multicast sources allowed to be tracked per group. Values 1 to 32000

max-num-sources

Syntax	max-num-sources <i>max-num-sources</i> no max-num-sources
Context	config>service>vprn>sub-if>grp-if>sap>igmp-host-tracking
Description	This command specifies the maximum number of multicast sources allowed to be tracked per group. The no form of the command reverts to the default.
Default	no max-num-sources
Parameters	<i>max-num-sources</i> — Specifies the maximum number of multicast sources allowed to be tracked per group.
Values	1 to 1000

label-mode

Syntax	label-mode { vrf next-hop } no label-mode
Context	config>service>vprn
Description	This command controls the method by which service labels are allocated to routes exported by the VPRN as BGP-VPN routes. The vrf option selects service label per VRF mode while the next-hop option selects service label per next-hop mode. The no form of the command sets the mode to the default mode of service label per VRF.
Default	no label-mode
Parameters	vrf — Selects service label per VRF mode. next-hop — Selects service label per next-hop mode.

maximum-ipv6-routes

Syntax	maximum-ipv6-routes <i>number</i> [log-only] [threshold <i>percentage</i>] no maximum-ipv6-routes
Context	config>service>vprn
Description	This command specifies the maximum number of remote IPv6 routes that can be held within a VPN routing/ forwarding (VRF) context. The local , host , static and aggregate routes are not counted.

The VPRN service ID must be in a shutdown state in order to modify maximum-routes command parameters.

If the **log-only** parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then the offending RIP peer (if applicable) is brought down (but the VPRN instance remains up). BGP peering will remain up but the exceeding BGP routes will not be added to the VRF.

The maximum route threshold can dynamically change to increase the number of supported routes even when the maximum has already been reached. Protocols will resubmit their routes which were initially rejected.

The **no** form of the command disables any limit on the number of routes within a VRF context. Issue the **no** form of the command only when the VPRN instance is shutdown.

Default	0 or disabled — The threshold will not be raised.
Parameters	<i>number</i> — An integer that specifies the maximum number of routes to be held in a VRF context.
Values	1 to 2147483647
	log-only — Specifies that if the maximum limit is reached, only log the event. log-only does not disable the learning of new routes.
	threshold percentage — The percentage at which a warning log message and SNMP trap should be set. There are two warnings, the first is a mid-level warning at the threshold value set and the second is a high-level warning at level between the maximum number of routes and the mid-level rate ($(\text{mid} + \text{max}) / 2$).
Values	0 to 100

maximum-routes

Syntax	maximum-routes <i>number</i> [log-only] [threshold percentage] no maximum-routes
Context	config>service>vprn
Description	This command specifies the maximum number of remote routes that can be held within a VPN routing/ forwarding (VRF) context. The local , host , static and aggregate routes are not counted.

The VPRN service ID must be in a shutdown state in order to modify maximum-routes command parameters.

If the **log-only** parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then the offending RIP peer (if applicable) is brought down (but the VPRN instance remains up). BGP peering will remain up but the exceeding BGP routes will not be added to the VRF.

The maximum route threshold can dynamically change to increase the number of supported routes even when the maximum has already been reached. Protocols will resubmit their routes which were initially rejected.

The **no** form of the command disables any limit on the number of routes within a VRF context. Issue the **no** form of the command only when the VPRN instance is shutdown.

Default	0 or disabled — The threshold will not be raised.
Parameters	<i>number</i> — An integer that specifies the maximum number of routes to be held in a VRF context.
Values	1 to 2147483647
	log-only — Specifies that if the maximum limit is reached, only log the event. log-only does not disable the learning of new routes.
	threshold percentage — The percentage at which a warning log message and SNMP trap should be set. There are two warnings, the first is a mid-level warning at the threshold value set and the second is a high-level warning at level between the maximum number of routes and the mid-level rate $([mid+max] / 2)$.
Values	0 to 100

multicast-info-policy

Syntax	multicast-info-policy <i>policy-name</i> no multicast-info-policy
Context	config>service>vprn
Description	This command configures multicast information policy.
Parameters	<i>policy-name</i> — Specifies the policy name.
Values	32 chars max

mc-maximum-routes

Syntax	mc-maximum-routes <i>number</i> [log-only] [threshold <i>threshold</i>]
Context	config>service>vprn
Description	This command specifies the maximum number of multicast routes that can be held within a VPN routing/forwarding (VRF) context. When this limit is reached, a log and SNMP trap are sent. If the log-only parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then no new joins will be processed.
	The no form of the command disables the limit of multicast routes within a VRF context. Issue the no form of the command only when the VPRN instance is shutdown.

Default	no mc-maximum-routes
Parameters	<i>number</i> — Specifies the maximum number of routes to be held in a VRF context.
Values	1 to 2147483647
	log-only — Specifies that if the maximum limit is reached, only log the event. log-only does not disable the learning of new routes.
	threshold <i>threshold</i> — The percentage at which a warning log message and SNMP trap should be sent.
Values	0 to 100
Default	10

network

Syntax	network
Context	config>service>vprn
Description	This command enters the context to configure network parameters for the VPRN service.

ingress

Syntax	ingress
Context	config>service>vprn>network
Description	This command enters the context to configure network ingress parameters for the VPRN service.

filter

Syntax	filter ip <i>ip-filter-id</i> filter ipv6 <i>ipv6-filter-id</i> no filter [ip <i>ip-filter-id</i>] [ipv6 <i>ipv6-filter-id</i>]
Context	config>service>vprn>network>ingress
Description	This command configures a network ingress filter for IPv4 or IPv6 traffic arriving over explicitly defined spokes or auto-bind network interfaces for the VPRN service. The no form of the command removes an IPv4, IPv6, or both filters.
Default	no filter

Parameters	<i>ip-filter-id</i> / <i>ipv6-filter-id</i> — Specifies an existing IP/IPv6 filter policy of a scope template.
Values	1 to 65535, <i>name</i> <i>name</i> : 64 characters maximum

ptp

Syntax	[no] ptp
Context	config>service>vprn
Description	This command enters the context to configure PTP parameters for the VPRN service.

peer-limit

Syntax	peer-limit <i>limit</i> no peer-limit
Context	config>service>vprn>ptp
Description	This command specifies an upper limit to the number of discovered peers permitted within the routing instance. This can be used to ensure that a routing instance does not consume all the possible discovered peers and blocking discovered peers in other routing instances. If it is desired to reserve a fixed number of discovered peers per router instance, then all router instances supporting PTP should have values specified with this command and the sum of all the peer-limit values must not exceed the maximum number of discovered peers supported by the system. If the user attempts to specify a peer-limit, and there are already more discovered peers in the routing instance than the new limit being specified, the configuration will not be accepted.
Default	no limit
Parameters	<i>limit</i> — Specifies the maximum number of discovered peers allowed in the routing instance. Values 0 to 50 Default 0 (The maximum number of discovered peers supported by the system).

peer

Syntax	peer <i>a.b.c.d</i> [create]
Context	config>system>ptp config>service>vprn>ptp

Description	This command configures a remote PTP peer. It provides the context to configure parameters for the remote PTP peer. Up to 20 remote PTP peers may be configured. The no form of the command deletes the specified peer. If the clock-type is ordinary slave or boundary, and PTP is no shutdown, the last peer cannot be deleted. This prevents the user from having PTP enabled without any peer configured and enabled. Peers are created within the routing instance associated with the context of this command. All configured PTP peers must use the same routing instance.
Parameters	<i>a.b.c.d</i> — The IP address of the remote peer.
Values	ipv4-address <i>a.b.c.d</i>

log-sync-interval

Syntax	log-sync-interval <i>log-interval</i>
Context	config>service>vprn>ptp>peer
Description	This command configures the message interval used for unicast event messages. It defines the message interval for both Sync and Delay_Resp messages that are requested during unicast negotiation to the specific peer. This controls the Sync and Delay_Resp message rate sent from remote peers to the local node. It does not affect the Sync or Delay_Resp packet rate that may be sent from the local node to remote peers. Remote peers may request a Sync or Delay_Resp packet rate anywhere within the acceptable grant range. The log-sync-interval cannot be changed unless the peer is shutdown. This command only applies to the 7450 ESS and 7750 SR.
Default	-6 (64 packets per second) for 1588-2008 or -6 (64 packets per second) for g8265dot1-2010 or -4 (16 packets per second) for g8275dot1-2014
Parameters	<i>log-interval</i> — Specifies the sync message interval, in log form.
Values	-6 to 0

local-priority

Syntax	local-priority <i>local-priority</i>
Context	config>service>vprn>ptp>peer

Description	This command configures the local priority used to choose between PTP masters in the best master clock algorithm (BMCA). This setting is relevant when the profile is set to either g8265dot1-2010 or g8275dot1-2014. The parameter is ignored when any other profile is selected. The value 1 is the highest priority and 255 is the lowest priority. The priority of a peer cannot be configured if the PTP profile is ieee1588-2008. For g8265dot1-2010, this parameter configures the priority used to choose between master clocks with the same quality (see G.8265.1 for more information). For g8275dot1-2014, this parameter sets the value of the localPriority associated with the Announce messages received from external clocks (ptp>peer or ptp>port), or the local clock (ptp). See G.8275.1 for more information.
Default	local-priority 128
Parameters	<i>local-priority</i> — Specifies the value of the local priority. Values 1 to 255

reassembly-group

Syntax	reassembly-group <i>nat-group-id</i> no reassembly-group
Context	config>router config>service>vprn
Description	This command associate reassembly-group consisting of multiple ISAs with the routing context in which the application requiring reassembly service resides.
Default	no reassembly-group
Parameters	<i>nat-group-id</i> — nat-group id; the nat-group contains up to 10 active ISAs. <i>asn:number</i> — The ASN is a 2-byte value less than or equal to 65535. The assigned number can be any 32-bit unsigned integer value.

route-distinguisher

Syntax	route-distinguisher [<i>ip-address:number</i> <i>asn:number</i>] route-distinguisher auto-rd no route-distinguisher
Context	config>service>vprn

Description	This command sets the identifier attached to routes the VPN belongs to. Each routing instance must have a unique (within the carrier's domain) route distinguisher associated with it. A route distinguisher must be defined for a VPRN to be operationally active. Alternatively, the auto-rd option allows the system to automatically generate a Route Distinguisher (RD) based on the bgp-auto-rd-range command configured at the service level.
Default	no route-distinguisher
Parameters	route distinguisher — The route distinguisher is a 6-byte value that can be specified in one of the following formats: <i>ip-address:number</i> — Specifies the IP address in dotted decimal notation. The assigned number must not be greater than 65535. <i>asn:number</i> — The ASN is a 2-byte value less than or equal to 65535. The assigned number can be any 32-bit unsigned integer value. auto-rd — The system will generate an RD for the service according to the IP address and range configured in the bgp-auto-rd-range command.

router-id

Syntax	router-id <i>ip-address</i> no router-id
Context	config>service>vprn config>service>vprn>ospf config>service>vprn>bgp
Description	This command sets the router ID for a specific VPRN context. When configuring the router ID in the base instance of OSPF it 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. If neither the router ID nor system interface are defined, the router ID from the base router context is inherited. This is a required command when configuring multiple instances and the instance being configured is not the base instance. 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.

It is possible to configure an SR OS to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP as there is no mechanism to derive the router ID from an IPv6 system interface address.

The **no** form of the command removes the router ID definition from the given VPRN context.

Default	no router-id
Parameters	<i>ip-address</i> — The IP address must be given in dotted decimal notation.

sgt-qos

Syntax	sgt-qos
Context	config>service>vprn
Description	This command configures DSCP/dot1p remarking for self-generated traffic.

application

Syntax	application <i>dscp-app-name</i> dscp { <i>dscp-value</i> <i>dscp-name</i> } application <i>dot1p-app-name</i> dot1p <i>dot1p-priority</i> no application { <i>dscp-app-name</i> <i>dot1p-app-name</i> }
Context	config>service>vprn>sgt-qos
Description	This command configures DSCP/dot1p remarking for self-generated traffic. When an application is configured using this command, then the specified DSCP name/value is used for all packets generated by this application within the router instance it is configured.

Using the value configured in this command:

- Sets the DSCP bits in the IP packet.
- Maps to the FC. This value will be signaled from the CPM to the egress forwarding complex.
- Based on this signaled FC the egress forwarding complex QoS policy sets the IEEE 802.1p and MPLS EXP bits.
- The DSCP value in the egress IP header will be as configured in this command. The egress QoS policy will not overwrite this value.

Only one DSCP name/value can be configured per application, if multiple entries are configured then the subsequent entry overrides the previous configured entry.

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

Parameters	<i>dscp-app-name</i> — Specifies the DSCP application name. Values bgp, bmp, dhcp, diameter, dns, ftp, gtp, icmp, igmp, l2tp, ldp, mld, msdp, ndis, ntp, ospf, pcep, pim, ptp, radius, rip, rsvp, snmp, snmp-notification, srrp, ssh, syslog, tacplus, telnet, tftp, traceroute, vrrp <i>dscp-value</i> — Specifies the DSCP value. Values 0 to 63 <i>dscp-name</i> — Specifies the DSCP name. Values none, be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63 <i>dot1p-priority</i> — Specifies the dot1p priority. Values none, 0 to 7 <i>dot1p-app-name</i> — Specifies the dot1p application name. Values arp, isis, pppoe

dscp

Syntax	dscp <i>dscp-name</i> fc <i>fc-name</i> no dscp <i>dscp-name</i>
Context	config>service>vprn>sgt-qos
Description	This command creates a mapping between the DiffServ Code Point (DSCP) of the self-generated traffic and the forwarding class.

Self generated traffic that matches the specified DSCP will be assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all sixty-four DiffServ code points to the forwarding class. For undefined code points, packets are assigned to the forwarding class specified under the default-action command.

All DSCP names that defines a DSCP value must be explicitly defined.

The **no** form of this command removes the DiffServ code point to forwarding class association. The default-action then applies to that code point value.

Parameters	<i>dscp-name</i> — Specifies the DSCP name.
Values	be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63
	fc <i>fc-name</i> — Specifies the forwarding class name.
Values	be, l2, af, l1, h2, ef, h1, nc

single-sfm-overload

Syntax	single-sfm-overload [holdoff-time <i>holdoff-time</i>] no single-sfm-overload
Context	config>service>vprn
Description	This command configures OSPF, OSPFv3 and IS-IS to set overload when the router has fewer than the full set of SFMs functioning, which reduces forwarding capacity. Setting overload enables a router to still participate in exchanging routing information, but routes all traffic away from it. The conditions to set overload are as follows: • 7950 XRS-20, 7750 SR-12/SR-7/SR-c12, and 7450 ESS-12/ESS-7 platforms: if an SF/CPMs fails, the protocol will set the overload • 7950-40 XRS and 7750 SR-12e platforms: if two SFMs fail (a connected pair on the XRS-40) the protocol will set the overload The no form of this command configures the router to not set overload if an SFM fails.
Default	no single-sfm-overload
Parameters	<i>holdoff-time</i> — Specifies the delay between detecting SFM failures and setting overload.
Values	1 to 600 seconds
Default	0 seconds

snmp

Syntax	snmp
Context	config>service>vprn
Description	This command enters the context to configure SNMP parameters for this VPRN.

access

Syntax	[no] access
Context	config>service>vprn>snmp
Description	This command enables/disables SNMP access on the VPRN interface. This command allows SNMP queries destined to the VPRN interface IP addresses for this VPRN (including VPRN interfaces that are bound to R-VPLS services) to be processed by the SNMP agent on the router. SNMP queries that arrive on VPRN interfaces but are destined to IP addresses in the Base routing context that can be accessed in the VPRN (for example, the router system address via grt leaking) do not require snmp-access to be enabled but do require allow-local-management to be enabled.

Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR System Management Guide* for detailed information about SNMP.

community

Syntax	community <i>community-name</i> [hash hash2] [access-permissions] [version <i>SNMP-version</i>] [src-access-list <i>list-name</i>] no community [<i>community-name</i>]
Context	config>service>vprn>snmp
Description	This command sets the SNMP community name(s) to be used with the associated VPRN instance. These VPRN community names are used to associate SNMP v1/v2c requests with a particular vprn context and to return a reply that contains VPRN-specific data or limit SNMP access to data in a specific VPRN instance.

VPRN snmp communities configured with an access permission of 'r' are automatically associated with the default access group "snmp-vprn-ro" and the "vprn-view" view (read only). VPRN snmp communities configured with an access permission of 'rw' are automatically associated with the default access group "snmp-vprn" and the "vprn-view" view (read/write).

The community in an SNMP v1/v2 request determines the SNMP context (i.e., the vprn# for accessing SNMP tables) and not the VPRN of the incoming interface on which the request was received. When an SNMP request arrives on VPRN 5 interface "ringo" with a destination IP address equal to the "ringo" interface, but the community in the SNMP request is the community configured against VPRN 101, then the SNMP request will be processed using the VPRN 101 context. (the response will contain information about VPRN 101). It is recommended to avoid using a simple series of vprn snmp-community values that are similar to each other (for example, avoid my-vprncomm-1, my-vprn-comm-2, etc).

The **no** form of the command removes the SNMP community name from the given VPRN context.

Parameters	<i>community-name</i> — Specifies the SNMP v1/v2c community name. This is a secret/confidential key used to access SNMP and specify a context (base vs vprn1 vs vprn2). hash — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified. hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the hash2 encrypted variable cannot be copied and pasted. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified. version <i>SNMP-version</i> — Specifies the SNMP version. Values v1, v2c, both <i>access-permissions</i> — Specifies the access rights to MIB objects. Values r — Grants only read access to MIB objects. Creates an association of the community-name with the snmp-vprn-ro access group. rw — Grants read and write access to MIB objects. Creates an association of the community-name with the snmp-vprn access group. <i>list-name</i> — Configures the community to reference a specific src-access-list (created under configure system security snmp), which will be used to validate the source IP address of all received SNMP requests that use this community. Multiple community (vprn or base router) and usm-community instances can reference the same src-access-list.

source-address

Syntax	source-address
Context	config>service>vprn
Description	This command enters the context to specify the source address and application that should be used in all unsolicited packets.

application

Syntax	application <i>app</i> [<i>ip-int-name</i> <i>ip-address</i>] no application <i>app</i>
Context	config>service>vprn>source-address
Description	This command specifies the source address and application.

Parameters	<i>app</i> — Specifies the application name.
Values	cflowd, dns, ftp, ntp, ping, ptp, radius, snmptrap, sntp, ssh, syslog, tacplus, telnet, traceroute, mcreporter, icmp-error
	<i>ip-int-name</i> <i>ip-address</i> — Specifies the name of the IP interface or IP address. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

application6

Syntax	application6 <i>app</i> <i>ipv6-address</i>
Context	config>service>vprn>source-address
Description	This command specifies the IPv6 source address and application.
Parameters	<i>app</i> — Specifies the application name.
Values	cflowd, dns, ftp, ntp, ping, radius, snmptrap, syslog, tacplus, telnet, traceroute, icmp6-error
	<i>ipv6-address</i> — Specifies the IPv6 address.

static-route-entry

Syntax	static-route-entry { <i>ip-prefix</i> / <i>prefix-length</i> } [mcast] no static-route-entry { <i>ip-prefix</i> / <i>prefix-length</i> } [mcast]
Context	config>service>vprn
Description	This command creates a static route entry for both the network and access routes. A prefix and netmask must be specified. Once the static route context for the specified prefix and netmask has been created, additional parameters associated with the static route(s) may be specified through the inclusion of additional static-route parameter commands. The no form of the command deletes the static route entry. If a static route needs to be removed when multiple static routes exist to the same destination, then as many parameters to uniquely identify the static route must be entered. IPv6 static routes are not supported on the 7450 ESS except in mixed mode.
Default	No static routes are defined.
Parameters	<i>ip-prefix/prefix-length</i> — The destination address of the static route.
Values	The following values apply to the 7750 SR and 7950 XRS:

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 to 32
ipv6-prefix	x::x::x::x::x::x (eight 16-bit pieces) x::x::x::x::x::d.d.d.d x: [0 to FFFF]H d: [0 to 255]D
ipv6-prefix-length	0 to 128

Values The following values apply to the 7450 ESS:

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 to 32

mcast — Specifies that the associated static route should be populated in the associated VPRN multicast route table.

next-hop

Syntax	next-hop { <i>ip-address</i> <i>ip-int-name</i> <i>ipv6 address</i> }						
Context	config>service>vprn>static-route-entry						
Description	This command specifies the directly connected next hop IP address or interface used to reach the destination. If the next hop is over an unnumbered interface or a point-to-point interface, the <i>ip-int-name</i> of the unnumbered or point-to-point interface (on this node) can be configured. If the next hop is over an unnumbered interface in the 7450 ESS router, the <i>ip-int-name</i> of the unnumbered interface (on this node) can be configured. The configured <i>ip-address</i> can be either on the network side or the access side on this node. The address must be associated with a network directly connected to a network configured on this node.						
Default	no next-hop						
Parameters	<i>ip-int-name</i> , <i>ipv4-address</i> , <i>ipv6-address</i> — the IP-INT, IPv4, and IPv6 addresses						
Values	The following values apply to the 7750 SR and 7950 XRS: <table> <tr> <td>ip-int-name</td><td>32 characters max</td></tr> <tr> <td>ipv4-address</td><td>a.b.c.d</td></tr> <tr> <td>ipv6-address</td><td>x::x::x::x::x::x-[interface] x::x::x::x::x::d.d.d.d[-interface] x: [0 to FFFF]H</td></tr> </table>	ip-int-name	32 characters max	ipv4-address	a.b.c.d	ipv6-address	x::x::x::x::x::x-[interface] x::x::x::x::x::d.d.d.d[-interface] x: [0 to FFFF]H
ip-int-name	32 characters max						
ipv4-address	a.b.c.d						
ipv6-address	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

IPv6 static routes are not supported on the 7450 ESS except in mixed mode.

indirect

Syntax	[no] indirect <i>ip-address</i>				
Context	config>service>vprn>static-route-entry				
Description	This command specifies that the route is indirect and specifies the next hop IP address used to reach the destination. The configured <i>ip-address</i> is not directly connected to a network configured on this node. The destination can be reached via multiple paths. The indirect address can only be resolved from dynamic routing protocol. Another static route cannot be used to resolve the indirect address. The <i>ip-address</i> configured here can be either on the network side or the access side and is typically at least one hop away from this node.				
Default	no indirect				
Parameters	<i>ip-address</i> — The IP address of the IP interface.				
	Values				
	<table> <tr> <td>ipv4-address</td><td>a.b.c.d</td></tr> <tr> <td>ipv6-address</td><td>x:x:x:x:x:x-x-[interface]</td></tr> </table>	ipv4-address	a.b.c.d	ipv6-address	x:x:x:x:x:x-x-[interface]
ipv4-address	a.b.c.d				
ipv6-address	x:x:x:x:x:x-x-[interface]				

black-hole

Syntax	[no] black-hole
Context	config>service>vprn>static-route-entry
Description	This command specifies that the route is a black hole route. If the destination address on a packet matches this static route, it will be silently discarded.
Default	no black-hole

community

Syntax **community** *comm-id* [*comm-id*]

no community

Context	config>service>vprn>static-route-entry
Description	This command associates a list of up to 12 BGP communities (any mix of standard, extended, and large communities) with the static route. These communities can be matched in route policies and are automatically added to BGP routes that are created from the static route. The communities specified at this level of the static route causes communities configured under the next-hop, black-hole, and indirect contexts of the static route to be ignored. The no form of this command removes the association.
Default	no community
Parameters	<i>comm-id</i> — Specifies a BGP community value, up to 72 characters. Values [<i>as-num:comm-val</i> <i>well-known-comm</i> <i>ext-comm</i> <i>large-comm</i>] where: • <i>as-num</i> — 0 to 65535 • <i>comm-val</i> — 0 to 65535 • <i>well-known-comm</i> — null no-export no-export-subconfed no-advertise llgr-stale no-llgr blackhole • <i>ext-comm</i> — the extended community, defined as one of the following: – {<i>target</i> <i>origin</i>}:<i>ip-address:comm-val</i> – {<i>target</i> <i>origin</i>}:<i>asnum:ext-comm-val</i> – {<i>target</i> <i>origin</i>}:<i>ext-asnum:comm-val</i> – bandwidth:<i>asnum:val-in-mbps</i> – ext:4300:<i>ovstate</i> – ext:<i>value1:value2</i> – color:<i>co-bits:color-value</i> where: – <i>target</i> — route target – <i>origin</i> — route origin – <i>ip-address</i> — a.b.c.d – <i>ext-comm-val</i> — 0 to 4294967295 – <i>ext-asnum</i> — 0 to 4294967295 – <i>val-in-mbps</i> — 0 to 16777215 – <i>ovstate</i> — 0, 1, or 2 (0 for valid, 1 for not found, 2 for invalid) – <i>value1</i> — 0000 to FFFF – <i>value2</i> — 0 to FFFFFFFFFFFFFFFF – <i>co-bits</i> — 00, 01, 10 or 11 – <i>color-value</i> — 0 to 4294967295

• *large-comm* — *asn-or-ex:val-or-ex:val-or-ex*

grt

Syntax	[no] grt
Context	config>service>vprn>static-route-entry
Description	This command creates a static route in a VPRN service context that points to the global routing context (base router). This is primarily used to allow traffic that ingress through a VPRN service to be routed out of the global routing context. This next-hop type cannot be used in conjunction with any other next-hop types.
Default	no grt

ipsec-tunnel

Syntax	[no] ipsec-tunnel <i>name</i>
Context	config>service>vprn>static-route-entry
Description	This command creates a static route in a VPRN service context that points to the global routing context (base router). This is primarily used to allow traffic that ingress through a VPRN service to be routed out of the global routing context. This next-hop type cannot be used in conjunction with any other next-hop types.
Default	no ipsec-tunnel
Parameters	<i>name</i> — IPsec tunnel name; maximum length up to 32 characters.

tag

Syntax	tag <i>tag</i> no tag
Context	config>service>vprn>static-route-entry
Description	This command associates a 4-byte route-tag with the static route. The tag value can be used in route policies to control distribution of the static route into other protocols. The tag specified at this level of the static route causes tag values configured under the next-hop, black-hole, and indirect contexts of the static route to be ignored. The no form of this command removes the tag association.

Default	no tag
Parameters	<i>tag</i> — Specifies an integer value.
Values	1 to 4294967295

bfd-enable

Syntax	[no] bfd-enable
Context	config>service>vprn>static-route-entry>next-hop
Description	This command associates the static route state to a BFD session between the local system and the configured nexthop. The remote end of the BFD session must also be configured to originate or accept the BFD session controlling the static route state. The no form of this command removes the association of the static route state to that of the BFD session.
Default	no bfd-enable

community

Syntax	community <i>comm-id</i> no community		
Context	config>service>vprn>static-route-entry>black-hole config>service>vprn>static-route-entry>indirect config>service>vprn>static-route-entry>next-hop		
Description	This command associates one BGP community (standard, extended or large) with a next-hop of the static route. This community can be matched in route policies and automatically added to BGP routes that are created from the static route. Any community specified in one of these contexts is overridden by any communities specified at the prefix level of the static route entry. The no form of this command removes the association.		
Default	no community		
Parameters	<i>comm-id</i> — Specifies a BGP community value, up to 72 characters. <table> <tr> <td>Values</td><td>[<i>as-num:comm-val</i> <i>well-known-comm</i> <i>ext-comm</i> <i>large-comm</i>]</td></tr> </table> where: • <i>as-num</i> — 0 to 65535 • <i>comm-val</i> — 0 to 65535	Values	[<i>as-num:comm-val</i> <i>well-known-comm</i> <i>ext-comm</i> <i>large-comm</i>]
Values	[<i>as-num:comm-val</i> <i>well-known-comm</i> <i>ext-comm</i> <i>large-comm</i>]		

- *well-known-comm* — **null** | **no-export** | **no-export-subconfed** | **no-advertise** | **llgr-stale** | **no-llgr** | **blackhole**
 - *ext-comm* — the extended community, defined as one of the following:
 - {*target* | *origin*}:*ip-address*:*comm-val*
 - {*target* | *origin*}:*asnum*:*ext-comm-val*
 - {*target* | *origin*}:*ext-asnum*:*comm-val*
 - **bandwidth**:*asnum*:*val-in-mbps*
 - **ext:4300**:*ovstate*
 - **ext**:*value1*:*value2*
 - **color**:*co-bits*:*color-value*
- where:
- *target* — route target
 - *origin* — route origin
 - *ip-address* — a.b.c.d
 - *ext-comm-val* — 0 to 4294967295
 - *ext-asnum* — 0 to 4294967295
 - *val-in-mbps* — 0 to 16777215
 - *ovstate* — 0, 1, or 2 (0 for valid, 1 for not found, 2 for invalid)
 - *value1* — 0000 to FFFF
 - *value2* — 0 to FFFFFFFFFFFFFFFF
 - *co-bits* — 00, 01, 10 or 11
 - *color-value* — 0 to 4294967295
- *large-comm* — *asn-or-ex:val-or-ex:val-or-ex*

community

Syntax	[no] community <i>comm-id</i>
Context	config>service>vprn>static-route-entry>ip-sec-tunnel
Description	This configuration option associates a BGP community with the static route. The community can be matched in route policies and is automatically added to BGP routes exported from the static route. The no form of this command removes the community association.
Default	no community
Parameters	<i>comm-id</i> — Specifies community IDs, up to 72 characters. Values [2 byte <i>asnumber:comm-val</i> <i>well-known-comm</i>] where:

- *2 byte as-number* — 0 to 65535
- *comm-val* — 0 to 65535
- *well-known-comm* — **no-export** | **no-export-subconfed** | **no-advertise**

cpe-check

Syntax	[no] cpe-check <i>cpe-ip-address</i>
Context	config>service>vprn>static-route-entry>indirect config>service>vprn>static-route-entry>next-hop
Description	This command enables CPE-check and specifies the IP address of the target CPE device. This option initiates a background ICMP ping test to the configured target IP address. The IP address can either be an IPv4 address for IPv4 static routes or an IPv6 address for IPv6 static routes. The target-ip-address cannot be in the same subnet as the static route subnet itself to avoid possible circular references. This option is mutually exclusive with BFD support on a given static route. The no form of this command disables the cpe-check option.
Default	no cpe-check
Parameters	<i>cpe-ip-address</i> — Specifies the IP address of the CPE device.

drop-count

Syntax	[no] drop-count <i>count</i>
Context	config>service>vprn>static-route-entry>indirect>cpe-check config>service>vprn>static-route-entry>next-hop>cpe-check
Description	This optional parameter specifies the number of consecutive ping-replies that must be missed to declare the CPE down and to deactivate the associated static route.
Default	drop-count 3
Parameters	<i>count</i> — An integer count value.
Values	1 to 255

interval

Syntax	[no] interval <i>seconds</i>
Context	config>service>vprn>static-route-entry>indirect>cpe-check

config>service>vprn>static-route-entry>next-hop>cpe-check

Description	This optional parameter specifies the interval between ICMP pings to the target IP address.
Default	interval 1
Parameters	<i>seconds</i> — An integer interval value.
Values	1 to 255

padding-size

Syntax	[no] padding-size <i>padding-size</i>
Context	config>service>vprn>static-route-entry>indirect>cpe-check config>service>vprn>static-route-entry>next-hop>cpe-check
Description	This optional parameter specifies the amount of padding to add to the ICMP packet in bytes. The parameter is only applicable when the cpe-check option is used with the associated static route.
Default	padding-size 56
Parameters	<i>padding-size</i> — An integer value.
Values	0 to 16384 bytes

log

Syntax	[no] log
Context	config>service>vprn>static-route-entry>indirect>cpe-check config>service>vprn>static-route-entry>next-hop>cpe-check
Description	This optional parameter enables the ability to log transitions between active and in-active based on the CPE connectivity check. Events will be sent to the system log, syslog and SNMP traps.
Default	no log

description

Syntax	[no] description <i>description-string</i>
Context	config>service>vprn>static-route-entry>indirect config>service>vprn>static-route-entry>next-hop config>service>vprn>static-route-entry>black-hole config>service>vprn>static-route-entry>grt

```
config>service>vprn>static-route-entry>ipsec-tunnel
```

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

The **no** form of the command removes the description string from the context

Default no description

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

destination-class

Syntax **[no] destination-class** *dest-index*

Context config>service>vprn>static-route-entry>indirect
config>service>vprn>static-route-entry>ipsec-tunnel

Description This command configures the policy accounting destination-class index to be used when incrementing accounting statistic for traffic matching the associated static route.

The **no** form of the command removes the associated destination-class from the associated static route nexthop.

Default no destination-class

Parameters *dest-index* — The destination index integer value.

Values 1 to 255

generate-icmp

Syntax **[no] generate-icmp**

Context config>service>vprn>static-route-entry>black-hole

Description This optional command causes the ICMP unreachable messages to be sent when received packets match the associated static route. By default, the ICMP unreachable messages for those types of static routes are not generated.

This command can only be associated with a static route that has a black-hole next-hop

The **no** form of this command removes the black-hole next-hop from static route configuration.

Default no generate-icmp

forwarding-class

Syntax	[no] forwarding-class { <i>be</i> <i>l2</i> <i>af</i> <i>l1</i> <i>h2</i> <i>ef</i> <i>h1</i> <i>nc</i> }
Context	config>service>vprn>static-route-entry>indirect config>service>vprn>static-route-entry>ipsec-tunnel
Description	This command specifies the enqueueing forwarding class that should be associated with traffic matching the associate static route. If this parameter is not specified, the packet will use the forwarding-class association based on default classification or other QoS Policy associations.
Default	no forwarding-class
Parameters	<i>Forwarding class</i> — The forwarding class must be one of the pre-defined system forwarding classes.
Values	be, l2, af, l1, h2, ef, h1, nc

ldp-sync

Syntax	[no] ldp-sync
Context	config>service>vprn>static-route-entry>indirect
Description	This command extends the LDP synchronization feature to a static route. When an interface comes back up, it is possible that a preferred static route using the interface as next-hop for a given prefix is enabled before the LDP adjacency to the peer LSR comes up on this interface. In this case, traffic on an SDP that uses the static route for the far-end address would be black-holed until the LDP session comes up and the FECs exchanged. This option when enabled delays the activation of the static route until the LDP session comes up over the interface and the ldp-sync-timer configured on that interface has expired
Default	no ldp-sync

load-balancing-weight

Syntax	load-balancing-weight <i>value</i> no load-balancing-weight
Context	config>service>vprn>static-route-entry>next-hop
Description	This command configures a weighted ECMP load-balancing weight for a static route next-hop. If all of the ECMP next-hops of a static route have a configured load-balancing-weight then packets matching the route are sprayed according to the relative weights. In other words, the next-hop interface with the largest load-balancing weight should receive the most forwarded traffic if weighted ECMP is applicable.

The **no** form of this command disables weighted ECMP for the interface and effectively disables weighted ECMP for the entire static route.

Parameters *value* — Specifies the cost metric value.

Values 0 to 4294967295

metric

Syntax **[no] metric** *metric-value*

Context config>service>vprn>static-route-entry>next-hop
config>service>vprn>static-route-entry>grt
config>service>vprn>static-route-entry>indirect
config>service>vprn>static-route-entry>ipsec-tunnel

Description This command specifies the cost metric for the static route, expressed as a decimal integer. This value is used when importing the static route into other protocols such as OSPF. When the metric is configured as 0 then the metric configured in OSPF, default-import-metric, applies. When modifying the metric of an existing static route, the preference will not change unless specified. This value is also used to determine which static route to install in the forwarding table.

If there are multiple static routes with the same preference but different metrics then the lower cost (metric) route will be installed.

The **no** form of this command returns the metric to the default value

Default metric 1

Parameters *metric-value* — Specifies the cost metric value.

Values 0 to 65535

preference

Syntax **[no] preference** *preference-value*

Context config>service>vprn>static-route-entry>grt
config>service>vprn>static-route-entry>indirect
config>service>vprn>static-route-entry>ipsec-tunnel
config>service>vprn>static-route-entry>black-hole

Description This command specifies the route preference to be assigned to the associated static route. The lower the preference value the more preferred the route is considered.

[Table 38](#) shows the default route preference based on the route source.

Table 38 **Default Route Preference**

Label	Preference	Configurable
Direct attached	0	No
Static route	5	Yes
OSPF Internal routes	10	Yes
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

The **no** form of this command returns the associated static route preference to its default value.

Default preference 5

Parameters *preference-value* — Specifies the route preference value.

Values 1 to 255

prefix-list

Syntax **[no] prefix-list** *name* {**all** | **none** | **any**}

Context config>service>vprn>static-route-entry>indirect
config>service>vprn>static-route-entry>next-hop
config>service>vprn>static-route-entry>black-hole

Description This command associates a new constraint to the associated static route such that the static route is only active if **any**, **none**, or **all** of the routes in the prefix list are present and active in the route-table.

Default no prefix-list

Parameters *name* — Specifies the name of a currently configured prefix-list.

all — Specifies that the static route condition is met if all prefixes in the prefix-list must be present in the active static route.

none — Specifies that the static route condition is met if none of the prefixes in the named prefix-list can be present in the active static route.

any — Specifies that the static route condition is met if any prefixes in the prefix-list are present in the active static route.

priority

Syntax	[no] priority {low high}
Context	config>service>vprn>static-route-entry>indirect>forwarding-class config>service>vprn>static-route-entry>next-hop>forwarding-class
Description	This optional command associates an enqueueing priority with the static route. The options are either high or low, with low being the default. This parameter has the ability to affect the likelihood that a packet will be enqueued at SAP ingress in the face of ingress congestion. Once a packet is enqueued into an ingress buffer, the significance of this parameter is lost.
Default	priority low
Parameters	low — Setting the enqueueing parameter for a packet to low decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. Once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost. high — Setting the enqueueing parameter for a packet to high increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. Once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

shutdown

Syntax	[no] shutdown
Context	config>service>vprn>static-route-entry>black-hole config>service>vprn>static-route-entry>grt config>service>vprn>static-route-entry>indirect config>service>vprn>static-route-entry>ipsec-tunnel config>service>vprn>static-route-entry>next-hop
Description	This command causes the static route to be placed in an administratively down state and removed from the active route-table
Default	no shutdown

source-class

Syntax **[no] source-class** *source-index*

Context	config>service>vprn>static-route-entry>indirect config>service>vprn>static-route-entry>ipsec-tunnel config>service>vprn>static-route-entry>next-hop
Description	This command configures the policy accounting source-class index to be used when incrementing accounting statistic for traffic matching the associated static route. If source route policy accounting is enabled and a source-class index is configured, traffic with a source IP address matches the associated static route, the source accounting statistics for the specified class will be incremented. The no form of the command removes the associated destination-class from the associated static route nexthop.
Default	no source-class
Parameters	<i>source-index</i> — Specifies an integer value for the accounting source class index.
Values	1 to 255

tag

Syntax	[no] tag <i>tag-value</i>
Context	config>service>vprn>static-route-entry>indirect config>service>vprn>static-route-entry>ipsec-tunnel config>service>vprn>static-route-entry>next-hop
Description	This command adds a 32-bit integer tag to the associated static route. The tag value can be used in route policies to control distribution of the route into other protocols.
Default	no tag
Parameters	<i>tag-value</i> — Specifies an integer tag value.
Values	32 bit integer

validate-next-hop

Syntax	[no] validate-next-hop
Context	config>service>vprn>static-route-entry>next-hop
Description	This optional command tracks the state of the next hop in the IPv4 ARP cache or IPv6 Neighbor Cache. When the next hop is not reachable and is removed from the ARP or Neighbor Cache, the next hop will no longer be considered valid and the associated static route state removed from the active route-table.

When the next hop is reachable again and present in the ARP/Neighbor Cache, the static route will be considered valid and is subject to being placed into the active route-table.

Default no validate-next-hop

ttl-propagate

Syntax **ttl-propagate**

Context config>service>vprn

Description This command enters the context to configure TTL propagation for transit and locally generated packets in a given VPRN routing context.

local

Syntax **local** [**inherit** | **all** | **vc-only** | **none**]

Context config>service>vprn>ttl-propagate

Description This command overrides the global configuration of the TTL propagation for locally generated packets which are forwarded over a MPLS LSPs in a given VPRN service context.

The global configuration is performed under config>router>ttl-propagate>vprn-local.

The default behavior for a given VPRN instance is to inherit the global configuration for the same command. The user can explicitly set the default behavior by configuring the inherit value

Default local inherit

Parameters *inherit* — specifies the TTL propagation behavior is inherited from the global configuration under config>router>ttl-propagate>vprn-local.

none — specifies the TTL of the IP packet is not propagated into the VC label or labels in the transport label stack.

vc-only — specifies the TTL of the IP packet is propagated into the VC label and not into the labels in the transport label stack.

all — specifies the TTL of the IP packet is propagated into the VC label and all labels in the transport label stack.

transit

Syntax **transit** [**inherit** | **all** | **vc-only** | **none**]

Context config>service>vprn>ttl-propagate

Description	This command overrides the global configuration of the TTL propagation for in transit packets which are forwarded over a MPLS LSPs in a given VPRN service context. The global configuration is performed under <code>config>router>ttl-propagate>vprn-transit</code>. The default behavior for a given VPRN instance is to inherit the global configuration for the same command. The user can explicitly set the default behavior by configuring the <code>inherit</code> value.
Default	transit inherit
Parameters	<i>inherit</i> — specifies the TTL propagation behavior is inherited from the global configuration under <code>config>router>ttl-propagate>vprn-transit</code>. <i>none</i> — specifies the TTL of the IP packet is not propagated into the VC label or labels in the transport label stack. <i>vc-only</i> — specifies the TTL of the IP packet is propagated into the VC label and not into the labels. in the transport label stack <i>all</i> — specifies the TTL of the IP packet is propagated into the VC label and all labels in the transport label stack.

type

Syntax	type [hub spoke subscriber-split-horizon] no type
Context	config>service>vprn>
Description	This command designates the type of VPRN instance being configured for hub and spoke topologies. Use the no form to reset to the default of a fully meshed VPRN.
Default	no type
Parameters	hub — Specifies a hub VPRN which allows all traffic from the hub SAPs to be routed to the destination directly, while all traffic from spoke VPRNs or network interfaces can only be routed to a hub SAP. spoke — Specifies a spoke VPRN which allows traffic from associated SAPs or spoke terminations to only be forwarded through routes learned from separate VPRN, which should be configured as a type Hub VPRN. subscriber-split-horizon — Controls the flow of traffic for wholesale subscriber applications.

vrf-export

Syntax	vrf-export <i>plcy-or-long-expr</i> [<i>plcy-or-expr</i> [<i>plcy-or-expr</i>]] no vrf-export

Context	config>service>vprn
Description	This command is used to specify route policies that control how routes are exported from the local VRF to other VRFs on the same or remote PE routers (via MP-BGP). Route policies are configured in the config>router>policy-options context. The vrf-export command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine final action to accept or reject the route. Only one of the 15 objects referenced by the vrf-export command can be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters while the remaining 14 objects have a maximum length of 64 characters each. When multiple vrf-export commands are issued, the last command entered overrides the previous command. Aggregate routes are not advertised via MP-BGP protocols to the other MP-BGP peers. The no form of the command removes all route policy names from the vrf-export list.
Default	no vrf-export
Parameters	<i>plcy-or-long-expr</i> — specifies the route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long). <i>plcy-or-expr</i> — specifies the route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long).

vrf-import

Syntax	vrf-import <i>plcy-or-long-expr</i> [<i>plcy-or-expr</i> [<i>plcy-or-expr</i>]] no vrf-import
Context	config>service>vprn
Description	This command is used to specify route policies that control how VPN-IP routes exported by other VRFs, on the same or remote PEs, are imported into the local VRF. Route policies are configured in the config>router>policy-options context. The vrf-import command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine final action to accept or reject the route. Only one of the 15 objects referenced by the vrf-import command is allowed to be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters while the remaining 14 objects have a maximum length of 64 characters each.

When multiple **vrf-import** commands are issued, the last command entered overrides the previous command.

The **no** form of the command removes all route policy names from the import list

Note that unless the preference value is changed by the policy, BGP-VPN routes imported with a vrf-import policy have the 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.

Default	no vrf-import
Parameters	<i>plcy-or-long-expr</i> — specifies the route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long). <i>plcy-or-expr</i> — specifies the route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long).

vrf-target

Syntax	vrf-target {ext-community export ext-community import ext-community} no vrf-target								
Context	config>service>vprn								
Description	This command facilitates a simplified method to configure the route target to be added to advertised routes or compared against received routes from other VRFs on the same or remote PE routers (via MP-BGP). BGP-VPN routes imported with a vrf-target statement will use the BGP preference value of 170 when imported from remote PE routers, or retain the protocol preference value of the exported route when imported from other VRFs in the same router. Specified vrf-import or vrf-export policies override the vrf-target policy. The no form of the command removes the vrf-target								
Default	no vrf-target								
Parameters	<i>ext-comm</i> — specifies the 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 <table> <tr> <td><ext-community></td><td>: target:{<ip-addr:comm-val> <2byte-asnumber:ext-comm-val> <4byte-asnumber:comm-val>}</td></tr> <tr> <td>ip-addr:</td><td>a.b.c.d</td></tr> <tr> <td>comm-val:</td><td>[0 to 65535]</td></tr> <tr> <td>2byte-asnumber:</td><td>[0 to 65535]</td></tr> </table>	<ext-community>	: 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-community>	: 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]

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.

weighted-ecmp

Syntax	weighted-ecmp no weighted-ecmp
Context	config>service>vprn
Description	This command enables weighted load-balancing for IS-IS ECMP routes in the VPRN instance. Weighted ECMP can be performed only when all the next hops are associated with the same neighbor and all of them are configured with (non-zero) load-balancing weights. The weighted ECMP support for IS-IS ECMP routes applies to both IPv4 and IPv6. The no form of the command restores regular ECMP spraying of packets to IS-IS route destinations.
Default	no weighted-ecmp

3.8.2.6 IPSec Configuration Commands

ipsec

Syntax	ipsec
Context	config>service>vprn>ipsec
Description	This command enters the context to configure IPSec policies.

security-policy

Syntax	security-policy <i>security-policy-id</i> [create] no security-policy <i>security-policy-id</i>
Context	config>service>vprn>ipsec

Description	This command configures a security policy to use for an IPSec tunnel.
Parameters	<i>security-policy-id</i> — Specifies a value to be assigned to a security policy.
	Values 1 to 8192
	create — Creates the security policy instance. The create keyword requirement can be enabled/disabled in the environment>create context.

entry

Syntax	entry <i>entry-id</i> [create] no entry <i>entry-id</i>
Context	config>service>vpn>ipsec>sec-plcy
Description	This command configures an IPSec security policy entry.
Parameters	<i>entry-id</i> — Specifies the IPSec security policy entry.
	Values 1 to 16
	create — Creates the security policy entry instance. The create keyword requirement can be enabled/disabled in the environment>create context.

local-ip

Syntax	local-ip { <i>ip-prefix/prefix-length</i> <i>ip-prefix netmask</i> any }
Context	config>service>vpn>ipsec>sec-plcy>entry
Description	This command configures the local (from the VPN) IP prefix/mask for the policy parameter entry. Only one entry is necessary to describe a potential flow. The local-ip and remote-ip commands can be defined only once. The system will evaluate the local IP as the source IP when traffic is examined in the direction of VPN to the tunnel and as the destination IP when traffic flows from the tunnel to the VPN. The remote IP will be evaluated as the source IP when traffic flows from the tunnel to the VPN when traffic flows from the VPN to the tunnel.
Parameters	<i>ip-prefix</i> — The destination address of the aggregate route in dotted decimal notation. Values a.b.c.d (host bits must be 0) prefix-length: 1 to 32 <i>netmask</i> — The subnet mask in dotted decimal notation. any — Specifies that it can be any address.

remote-ip

Syntax	remote-ip <i>ip-prefix/prefix-length</i> <i>ip-prefix netmask</i> any }				
Context	config>service>vpn>ipsec>sec-plcy>entry				
Description	This command configures the remote (from the tunnel) IP prefix/mask for the policy parameter entry. Only one entry is necessary to describe a potential flow. The local-ip and remote-ip commands can be defined only once. The system will evaluate the local IP as the source IP when traffic is examined in the direction of VPN to the tunnel and as the destination IP when traffic flows from the tunnel to the VPN. The remote IP will be evaluated as the source IP when traffic flows from the tunnel to the VPN when traffic flows from the VPN to the tunnel.				
Parameters	<i>ip-prefix</i> — The destination address of the aggregate route in dotted decimal notation. <table><tr><td>Values</td><td>a.b.c.d (host bits must be 0)</td></tr><tr><td></td><td>prefix-length: 1 to 32</td></tr></table> <i>netmask</i> — The subnet mask in dotted decimal notation. any — Specifies that it can be any address.	Values	a.b.c.d (host bits must be 0)		prefix-length: 1 to 32
Values	a.b.c.d (host bits must be 0)				
	prefix-length: 1 to 32				

interface

Syntax	interface <i>ip-int-name</i> [create] [tunnel] no interface <i>ip-int-name</i>
Context	config>service>vpn
Description	This command creates a logical IP routing interface for a Virtual Private Routed Network (VPRN). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface. The interface command, under the context of services, is used to create and maintain IP routing interfaces within VPRN service IDs. The interface command can be executed in the context of an VPRN service ID. The IP interface created is associated with the service core network routing instance and default routing table. The typical use for IP interfaces created in this manner is for subscriber internet access. Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for config router interface and config service vpn interface. Interface names must not be in the dotted decimal notation of an IP address. For example, the name "1.1.1.1" is not allowed, but "int-1.1.1.1" is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

The available IP address space for local subnets and routes is controlled with the **config router service-prefix** command. The **service-prefix** command administers the allowed subnets that can be defined on service IP interfaces. It also controls the prefixes that may be learned or statically defined with the service IP interface as the egress interface. This allows segmenting the IP address space into **config router** and **config service** domains.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All VPRN IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

The no form of this command removes IP the interface and all the associated configuration. The interface must be administratively shutdown before issuing the **no interface** command.

For VPRN services, the IP interface must be shutdown before the SAP on that interface may be removed. VPRN services do not have the **shutdown** command in the SAP CLI context. VPRN service SAPs rely on the interface status to enable and disable them.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

Values 1 to 32 characters maximum

tunnel — Specifies that the interface is configured as tunnel interface, which could be used to terminate IPsec or GRE tunnels in the private service.

create — Creates the IPsec interface instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

address

Syntax **[no] address** {*ip-address/mask* | *ip-address netmasks*}

Context config>service>vprn>if

Description This command assigns an IP address/IP subnet to the interface

Parameters *ip-address* — Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

mask — The subnet mask in dotted decimal notation. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. A mask of 255.255.255.255 is reserved for system IP addresses.

netmask — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.

ip-mtu

Syntax	ip-mtu <i>octets</i> no ip-mtu
Context	config>service>vprn>if config>service>vprn>if>sap>ip-tunnel config>service>vprn>sap>ipsec-tunnel
Description	This command configures the IP maximum transmit unit (packet) for this interface. The no form of the command returns the default value.
Default	no ip-mtu
Parameters	<i>octets</i> — Specifies the MTU size for this interface. Values 512 to 9000

ipsec-tunnel

Syntax	ipsec-tunnel <i>ipsec-tunnel-name</i> [create] no ipsec-tunnel <i>ipsec-tunnel-name</i>
Context	config>service>vprn>if>sap
Description	This command specifies an IPsec tunnel name. An IPsec client sets up the encrypted tunnel across public network. The 7750 SR IPsec MDA acts as a concentrator gathering, and terminating these IPsec tunnels into an IES or VPRN service. This mechanism allows as service provider to offer a global VPRN service even if node of the VPRN are on an uncontrolled or insecure portion of the network.
Parameters	<i>ipsec-tunnel-name</i> — Specifies an IPsec tunnel name up to 32 characters in length. create — Creates the IPsec tunnel instance. The create keyword requirement can be enabled/disabled in the environment>create context.

bfd-designate

Syntax	[no] bfd-designate
Context	config>service>vprn>if>sap>ipsec-tunnel
Description	This command specifies whether this IPsec tunnel is the BFD designated tunnel.

bfd-enable

Syntax	bfd-enable service <i>service-id</i> interface <i>interface-name</i> dst-ip <i>ip-address</i> bfd-enable service-name <i>service-name</i> interface <i>interface-name</i> dst-ip <i>ip-address</i> no bfd-enable
Context	config>service>vprn>if>sap>ipsec-tunnel
Description	This command assign a BFD session provide heart-beat mechanism for given IPsec tunnel. There can be only one BFD session assigned to any given IPsec tunnel, but there can be multiple IPsec tunnels using same BFD session. BFD control the state of the associated tunnel, if BFD session goes down, system will also bring down the associated non-designated IPsec tunnel.
Parameters	<i>service-id</i> — Specifies where the service-id that the BFD session resides. This variant of the command is only supported in 'classic' configuration-mode (configure system management-interface configuration-mode classic). The bfd-enable service-name <i>service-name</i> variant can be used in all configuration modes. Values {<i>id</i> <i>svc-name</i>} <i>id</i>: 1 to 2147483647 <i>svc-name</i>: up to 64 characters (<i>svc-name</i> is an alias for input only. The <i>svc-name</i> gets replaced with an id automatically by SR OS in the configuration). <i>interface-name</i> — Specifies the name of the interface used by the BFD session, up to 32 characters. <i>ip-address</i> — Specifies the destination address to be used for the BFD session. Values ipv4-address - a.b.c.d <i>service-name</i> — Identifies the service, up to 64 characters.

clear-df-bit

Syntax	[no] clear-df-bit
Context	config>service>interface>ies>sap config>service>vprn>if>sap>ipsec-tunnel config>service>vprn>if>sap>ip-tunnel config>service>ies>if>sap>ip-tunnel config>ipsec>tnl-temp
Description	This command specifies whether to clear the Do not Fragment (DF) bit in the outgoing packets in this tunnel.

dynamic-keying

Syntax	[no] dynamic-keying
Context	config>service>vprn>if>sap>ipsec-tunnel
Description	This command enables dynamic keying for the IPsec tunnel.

auto-establish

Syntax	[no] auto-establish
Context	config>service>vprn>if>sap>ipsec-tun>dyn
Description	This command specifies whether to attempt to establish a phase 1 exchange automatically. The no form of the command disables the automatic attempts to establish a phase 1 exchange.
Default	no auto-establish

local-id

Syntax	[no] local-id type {ipv4 <v4address> fqdn <fqdn-value>}
Context	config>service>vprn>if>sap>ipsec-tun>dynamic-keying
Description	This command specifies the local id of the 7750 SR used for IDi or IDr for IKEv2 tunnels. The local-id command can only be changed or removed when tunnel or gw is shutdown. The default value depends on the local-auth-method such as: • Psk:local tunnel ip address • Cert-auth: subject of the local certificate
Default	no local-id
Parameters	<i>type</i> — Specifies the type of local ID payload, it could be ipv4 address. <i>ipv4</i> — Specifies IPv4 as the local ID type. The default value is the local tunnel end-point address. <i>v4address</i> — Specifies an IPv4 address. A value must be configured. <i>fqdn</i> — Specifies FQDN as the local ID type. A value must be configured. <i>fqdn-value</i> — Specifies a FQDN value. A value must be configured.

transform

Syntax	transform <i>transform-id</i> [<i>transform-id</i>] no transform
Context	config>service>vprn>if>sap>ipsec-tun>dynamic-keying config>ipsec>tnl-temp
Description	This command associates the IPSec transform sets allowed for this tunnel. The transforms are listed in decreasing order of preference (the first one specified is the most preferred).
Parameters	<i>transform-id</i> — Specifies the value used for transforms for dynamic keying. A maximum of four transforms can be specified.
Values	1 to 2048

local-gateway-address

Syntax	local-gateway-address <i>ip-address</i> peer <i>ip-address</i> delivery-service <i>service-id</i> local-gateway-address <i>ip-address</i> peer <i>ip-address</i> delivery-service-name <i>service-name</i> no local-gateway-address
Context	config>service>vprn>if>sap>ipsec-tunnel
Description	This command specifies the local gateway address used for the tunnel and the address of the remote security gateway at the other end of the tunnel remote peer IP address to use.
Default	The base routing context is used if the delivery-router option is not specified.
Parameters	<i>ip-address</i> — IP address of the local end of the tunnel. <i>service-id</i> — The ID of the IES or VPRN (front-door) delivery service of this tunnel. Use this <i>service-id</i> to find the VPRN used for delivery. This variant of the command is only supported in 'classic' configuration-mode (configure system management-interface configuration-mode classic). The local-gateway-address delivery-service-name <i>service-name</i> variant can be used in all configuration modes.
Values	{ <i>id</i> <i>svc-name</i> } <i>id</i> : 1 to 2147483647 <i>svc-name</i> : up to 64 characters (<i>svc-name</i> is an alias for input only. The <i>svc-name</i> gets replaced with an id automatically by SR OS in the configuration).
	<i>service-name</i> — Identifies the service.
Values	1 to 64 characters

manual-keying

Syntax	[no] manual-keying
Context	config>service>vprn>if>sap>ipsec-tunnel
Description	This command configures Security Association (SA) for manual keying. When enabled, the command specifies whether this SA entry is created manually by the user or dynamically by the IPsec sub-system.

security-association

Syntax	security-association security-entry-id authentication-key authentication-key encryption-key encryption-key spi spi transform transform-id direction {inbound outbound} no security-association security-entry-id direction {inbound outbound}
Context	config>service>vprn>if>sap>ipsec-tun>manual-keying
Description	This command configures the information required for manual keying SA creation.
Parameters	security-entry-id — Specifies the ID of an SA entry. Values 1 to 16 encryption-key encryption-key — Specifies the key used for the encryption algorithm. Values none or 0x0 to 0xFFFFFFFF...(max 64 hex nibbles) authentication-key authentication-key — Specifies the key used for the authentication algorithm. Values none or 0x0 to 0xFFFFFFFF...(max 40 hex nibbles) spi spi — Specifies the SPI (Security Parameter Index) used to look up the instruction to verify and decrypt the incoming IPsec packets when the direction is inbound. When the direction is outbound, the SPI that will be used in the encoding of the outgoing packets. The remote node can use this SPI to lookup the instruction to verify and decrypt the packet. Values 256 to 16383 transform transform-id — Specifies the transform entry that will be used by this SA entry. This object should be specified for all the entries created which are manual SAs. If the value is dynamic, then this value is irrelevant and will be zero. Values 1 to 2048 direction {inbound outbound} — Specifies the direction of an IPsec tunnel.

replay-window

Syntax	replay-window {32 64 128 256 512}

no replay-window

Context	config>ipsec>tnl-temp config>service>vprn>if>sap>ipsec-tunnel
Description	This command specifies the size of the anti-replay window. The anti-replay window protocol secures IP against an entity that can inject messages in a message stream from a source to a destination computer on the Internet.
Parameters	{32 64 128 256 512} — Specifies the size of the SA anti-replay window.

security-policy

Syntax	security-policy <i>security-policy-id</i> no security-policy
Context	config>service>vprn>if>sap>ipsec-tunnel
Description	This command configures an IPSec security policy. The policy may then be associated with static IPsec tunnels defined in the same routing instance. With strict-match parameter enabled, when a CREATE_CHILD exchange request is received for a static IPsec tunnel, and this request is not a re-key request, then ISA matches the received TSi and TSr with the configured security policy. This can be a match only when a received TS (in TSi or TSr) address range matches exactly with the subnet in a security policy entry. If there is no match, then the setup fails, and TS_UNACCEPTABLE is sent. If there is a match, but there is an existing CHILD_SA for the matched security policy, then the setup fails, and NO_PROPOSAL_CHOSEN. If there is a match, and there is not CHILD_SA for the matched entry, then the subnet is sent in the matched security-policy entry as TSi and TSr, and the CHILD_SA is created.
Parameters	<i>security-policy-id</i> — Specifies the IPSec security policy entry that the tunnel will use. Values 1 to 8192 strict-match — Enables strict match of security-policy entry.

3.8.2.7 Log Commands

log

Syntax	log

Context	config>service>vprn config>service>vprn>log-id
Description	This command enters the context to configure event stream logging.

filter

Syntax	[no] filter <i>filter-id</i>
Context	config>service>vprn>log config>service>vprn>log>log-id
Description	This command creates a context for an event filter. An event filter specifies whether to forward or drop an event or trap based on the match criteria. Filters are configured in the filter <i>filter-id</i> context and then applied to a log in the log-id <i>log-id</i> context. Only events for the configured log source streams destined to the log ID where the filter is applied are filtered. Any changes made to an existing filter, using any of the sub-commands, are immediately applied to the destinations where the filter is applied. The no form of the command removes the filter association from log IDs which causes those logs to forward all events.
Default	No event filters are defined.
Parameters	<i>filter-id</i> — The filter ID uniquely identifies the filter.
Values	1 to 1000

default-action

Syntax	default-action {drop forward} no default-action
Context	config>service>vprn>log>filter
Description	The default action specifies the action that is applied to events when no action is specified in the event filter entries or when an event does not match the specified criteria. When multiple default-action commands are entered, the last command overwrites the previous command. The no form of the command reverts the default action to the default value (forward).
Default	default-action forward — The events which are not explicitly dropped by an event filter match are forwarded.

- Parameters**
- drop** — The events which are not explicitly forwarded by an event filter match are dropped.
 - forward** — The events which are not explicitly dropped by an event filter match are forwarded.

entry

- Syntax** `[no] entry entry-id`
- Context** `config>service>vprn>log>filter`
- Description** This command is used to create or edit an event filter entry. Multiple entries may be created using unique *entry-id* numbers. The TiMOS implementation exits the filter on the first match found and executes the action in accordance with the action command.
- Comparisons are performed in an ascending entry ID order. When entries are created, they should be arranged sequentially from the most explicit entry to the least explicit. Matching ceases when a packet matches an entry. The entry action is performed on the packet, either drop or forward. To be considered a match, the packet must meet all the conditions defined in the entry.
- An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword action for it to be considered complete. Entries without the action keyword will be considered incomplete and are rendered inactive.
- The **no** form of the command removes the specified entry from the event filter. Entries removed from the event filter are immediately removed from all log-id's where the filter is applied.
- Default** No event filter entries are defined. An entry must be explicitly configured.
- Parameters** *entry-id* — The entry ID uniquely identifies a set of match criteria corresponding action within a filter. Entry ID values should be configured in staggered increments so you can insert a new entry in an existing policy without renumbering the existing entries.
- Values** 1 to 999

action

- Syntax** `action {drop | forward}`
`no action`
- Context** `config>service>vprn>log>filter>entry`
- Description** This command specifies a drop or forward action associated with the filter entry. If neither **drop** nor **forward** is specified, the [default-action](#) will be used for traffic that conforms to the match criteria. This could be considered a No-Op filter entry used to explicitly exit a set of filter entries without modifying previous actions.

Multiple action statements entered will overwrite previous actions.

The **no** form of the command removes the specified [action](#) statement.

Default Action specified by the [default-action](#) command will apply.

Parameters **drop** — Specifies packets matching the entry criteria will be dropped.
forward — Specifies packets matching the entry criteria will be forwarded.

match

Syntax **[no] match**

Context config>service>vprn>log>filter>entry

Description This command creates context to enter/edit match criteria for a filter entry. When the match criteria is satisfied, the action associated with the entry is executed.

If more than one match parameter (within one match statement) is specified, then all the criteria must be satisfied (AND functional) before the action associated with the match is executed.

Use the **match** command to display a list of the valid applications.

Match context can consist of multiple match parameters (application, event-number, severity, subject), but multiple **match** statements cannot be entered per entry.

The **no** form of the command removes the match criteria for the *entry-id*.

Default no match

application

Syntax **application {eq | neq} application-id**
no application

Context config>service>vprn>log>filter>entry>match

Description This command adds an OS application as an event filter match criterion.

An OS application is the software entity that reports the event. Applications include IP, MPLS, OSPF, CLI, SERVICES and so on Only one application can be specified. The latest **application** command overwrites the previous command.

The **no** form of the command removes the application as a match criterion.

Default no application — no application match criterion is specified

Parameters **eq | neq** — The operator specifying the type of match.

Values

eq	equal to
neq	not equal to

application-id — The application name string.

Values

port, ppp, rip, route, policy, rsvp, security, snmp, stp, svcmgr,
system, user, vrrp, vrtr

message

Syntax **message {eq | neq} pattern pattern [regexp]**
no message

Context config>service>vprn>log>filter>entry>match

Description This command adds system messages as a match criterion.

The **no** form of the command removes messages as a match criterion.

Parameters **eq** — Determines if the matching criteria should be equal to the specified value.
neq — Determines if the matching criteria should not be equal to the specified value.
pattern pattern — Specifies a message up to 400 characters to be used in the match criteria.
regexp — Specifies the type of string comparison to use to determine if the log event matches the value of **message** command parameters. When the **regexp** keyword is not specified, the default matching algorithm used is a basic substring match.

number

Syntax **number {eq | neq | lt | lte | gt | gte} event-id**
no number

Context config>service>vprn>log>filter>entry>match

Description This command adds an SR OS application event number as a match criterion.

SR OS event numbers uniquely identify a specific logging event within an application.

Only one **number** command can be entered per event filter entry. The latest **number** command overwrites the previous command.

The **no** form of the command removes the event number as a match criterion.

Default	no event-number — No event ID match criterion is specified.
Parameters	eq neq lt lte gt gte — This operator specifies the type of match. Valid operators are listed below.

Values

Operator	Note
eq	equal to
neq	not equal to
lt	less than
lte	less than or equal to
gt	greater than
gte	greater than or equal to

event-id — Specifies the event ID, expressed as a decimal integer.

Values 1 to 4294967295

severity

Syntax	severity {eq neq lt lte gt gte} severity-level no severity
Context	config>service>vprn>log>filter>entry>match
Description	This command adds an event severity level as a match criterion. Only one severity command can be entered per event filter entry. The latest severity command overwrites the previous command. The no form of the command removes the severity match criterion.
Default	no severity
Parameters	eq neq lt lte gt gte — Specifies the type of match. Valid operators are listed below.

Values

Operator	Notes
eq	equal to
neq	not equal to
lt	less than
lte	less than or equal to
gt	greater than
gte	greater than or equal to

severity-name — The ITU severity level name. [Table 39](#) lists severity names and corresponding numbers per ITU standards M.3100 X.733 & X.21 severity levels.

Table 39 Severity Levels

Severity Number	Severity Name
1	cleared
2	indeterminate (info)
3	critical
4	major
5	minor
6	warning

Values cleared, intermediate, critical, major, minor, warning

subject

Syntax **subject** {**eq** | **neq**} *subject* [*regex*]
no subject

Context config>service>vprn>log>filter>entry>match

Description This command adds an event subject as a match criterion.

The subject is the entity for which the event is reported, such as a port. In this case the port-id string would be the subject. Only one **subject** command can be entered per event filter entry. The latest **subject** command overwrites the previous command.

The **no** form of the command removes the subject match criterion.

Default no subject

Parameters **eq** | **neq** — This operator specifies the type of match. Valid operators are listed below.

Values

Operator	Notes
eq	equal to
neq	not equal to

subject — A string used as the subject match criterion.

regex — Specifies the type of string comparison to use to determine if the log event matches the value of **subject** command parameters. When the **regex** keyword is specified, the string in the **subject** command is a regular expression string that will be matched against the subject string in the log event being filtered.

When **regex** keyword is not specified, the **subject** command string is matched exactly by the event filter.

log-id

Syntax	[no] log-id <i>log-id</i>		
Context	config>service>vprn>log		
Description	This command creates a context to configure destinations for event streams. The log-id context is used to direct events, alarms/traps, and debug information to respective destinations. A maximum of 30 logs can be configured. Before an event can be associated with this log-id, the from command identifying the source of the event must be configured. Only one destination can be specified for a <i>log-id</i>. The destination of an event stream can be an in-memory buffer, console, session, snmp-trap-group, syslog, or file. Use the event-control command to suppress the generation of events, alarms, and traps for all log destinations. An event filter policy can be applied in the log-id context to limit which events, alarms, and traps are sent to the specified log-id. Log-IDs 99 and 100 are created by the agent. Log-ID 99 captures all log messages. Log-ID 100 captures log messages with a severity level of major and above. Log-ID 99 provides valuable information for the admin-tech file. Removing or changing the log configuration may hinder debugging capabilities. It is strongly recommended not to alter the configuration for Log-ID 99. The no form of the command deletes the log destination ID from the configuration.		
Default	No log destinations are defined.		
Parameters	<i>log-id</i> — The log ID number, expressed as a decimal integer. <table><tr><td>Values</td><td>1 to 100</td></tr></table>	Values	1 to 100
Values	1 to 100		

to snmp

Syntax	to snmp [<i>size</i>]				
Context	config>service>vprn>log>log-id				
Description	This is one of the commands used to specify the log ID destination. This parameter is mandatory when configuring a log destination. This command instructs the alarms and traps to be directed to the snmp-trap-group associated with <i>log-id</i>. A local circular memory log is always maintained for SNMP notifications sent to the specified snmp-trap-group for the <i>log-id</i>. The source of the data stream must be specified in the from command prior to configuring the destination with the to command. The to command cannot be modified or re-entered. If the destination or maximum size of an SNMP or memory log needs to be modified, the log ID must be removed and then re-created.				
Parameters	<i>size</i> — The <i>size</i> parameter defines the number of events stored in this memory log. <table> <tr> <td>Default</td><td>100</td></tr> <tr> <td>Values</td><td>50 to 1024</td></tr> </table>	Default	100	Values	50 to 1024
Default	100				
Values	50 to 1024				

to syslog

Syntax	to syslog <i>syslog-id</i>		
Context	config>service>vprn>log>log-id		
Description	This is one of the commands used to specify the log ID destination. This parameter is mandatory when configuring a log destination. This command instructs the alarms and traps to be directed to a specified syslog. To remain consistent with the standards governing syslog, messages to syslog are truncated to 1k bytes. The source of the data stream must be specified in the from command prior to configuring the destination with the to command. The to command cannot be modified or re-entered. If the destination or maximum size of an SNMP or memory log needs to be modified, the log ID must be removed and then re-created.		
Parameters	<i>syslog-id</i> — Instructs the events selected for the log ID to be directed to the <i>syslog-id</i>. The characteristics of the <i>syslog-id</i> referenced here must have been defined in the config>log>syslog <i>syslog-id</i> context. <table> <tr> <td>Values</td><td>1 to 10</td></tr> </table>	Values	1 to 10
Values	1 to 10		

from

Syntax	from {[main] [change]} no from
Context	config>service>vprn>log>log-id
Description	This command selects the source stream to be sent to a log destination. One or more source streams must be specified. The source of the data stream must be identified using the from command before you can configure the destination using the to command. The from command can identify multiple source streams in a single statement (for example: from main change debug-trace). Only one from command may be entered for a single <i>log-id</i>. If multiple from commands are configured, then the last command entered overwrites the previous from command. The no form of the command removes all previously configured source streams.
Default	No source stream is configured.
Parameters	main — Instructs all events in the main event stream to be sent to the destination defined in the to command for this destination <i>log-id</i>. The main event stream contains the events that are not explicitly directed to any other event stream. To limit the events forwarded to the destination, configure filters using the filter command. change — Instructs all events in the user activity stream to be sent to the destination configured in the to command for this destination <i>log-id</i>. The change event stream contains all events that directly affect the configuration or operation of this node. To limit the events forwarded to the change stream destination, configure filters using the filter command.

time-format

Syntax	time-format {local utc}
Context	config>service>vprn>log>log-id
Description	This command specifies whether the time should be displayed in local or Coordinated Universal Time (UTC) format.
Default	time-format utc
Parameters	local — Specifies that timestamps are written in the system's local time. utc — Specifies that timestamps are written using the UTC value. This was formerly called Greenwich Mean Time (GMT) and Zulu time.

syslog

Syntax	[no] syslog <i>syslog-id</i>
Context	config>service>vprn>log
Description	This command creates the context to configure a syslog target host that is capable of receiving selected syslog messages from this network element. A valid <i>syslog-id</i> must have the target syslog host address configured. A maximum of 10 syslog-ids can be configured. No log events are sent to a syslog target address until the syslog-id has been configured as the log destination (to) in the log-id node. The syslog ID configured in the configure/service/vprn context has a local VPRN scope and only needs to be unique within the specific VPRN instance. The same ID can be reused under a different VPRN service or in the global log context under config>log.
Default	No syslog IDs are defined.
Parameters	<i>syslog-id</i> — Specifies the syslog ID number for the syslog destination, expressed as a decimal integer. Values 1 to 10

address

Syntax	address <i>ip-address</i> no address
Context	config>service>vprn>log>syslog
Description	This command adds the syslog target host IP address to/from a syslog ID. The <i>ip-address</i> parameter is mandatory. If no address is configured, syslog data cannot be forwarded to the syslog target host. Only one address can be associated with a <i>syslog-id</i>. If multiple addresses are entered, the last address entered overwrites the previous address. The same syslog target host can be used by multiple log IDs. The no form of the command removes the syslog target host IP address.
Default	no address

Parameters *ip-address* — Specifies the IP address of the syslog target host in dotted decimal notation.

Values

ipv4-address	a.b.c.d
ipv6-address	x:x:x:x:x:x:x[-interface] x:x:x:x:x:x.d.d.d[-interface] x: [0 to FFFF]H d: [0 to 255]D interface: 32 characters maximum, mandatory for link local addresses

The ipv6-address applies to the 7750 SR.

facility

Syntax **facility** *syslog-facility*
no facility

Context config>service>vprn>log>syslog

Description This command configures the facility code for messages sent to the syslog target host.

Multiple syslog IDs can be created with the same target host but each syslog ID can only have one facility code. If multiple facility codes are entered, the last *facility-code* entered overwrites the previous facility-code.

If multiple facilities need to be generated for a single syslog target host, then multiple **log-id** entries must be created, each with its own filter criteria to select the events to be sent to the syslog target host with a given facility code.

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

Default local7 — Syslog entries are sent with the local7 facility code.

Parameters *syslog-facility* — The syslog facility name represents a specific numeric facility code. The code should be entered in accordance with the syslog RFC. However, the software does not validate if the facility code configured is appropriate for the event type being sent to the syslog target host.

Values

kernel, user, mail, systemd, auth, syslogd, printer, netnews, uucp, cron, authpriv, ftp, ntp, logaudit, logalert, cron2, local0, local1, local2, local3, local4, local5, local6, local7

Valid responses per RFC3164, *The BSD syslog Protocol*, are listed in [Table 40](#).

Table 40 Syslog Facility Codes

Numerical Code	Facility Code
0	kernel
1	user
2	mail
3	systemd
4	auth
5	syslogd
6	printer
7	net-news
8	uucp
9	cron
10	auth-priv
11	ftp
12	ntp
13	log-audit
14	log-alert
15	cron2
16	local0
17	local1
18	local2
19	local3
20	local4
21	local5
22	local6
23	local7

Values: 0 to 23

log-prefix

Syntax	log-prefix <i>log-prefix-string</i> no log-prefix
Context	config>service>vprn>log>syslog
Description	This command adds the string prepended to every syslog message sent to the syslog host. RFC3164, <i>The BSD syslog Protocol</i>, allows an alphanumeric string (tag) to be prepended to the content of every log message sent to the syslog host. This alphanumeric string can, for example, be used to identify the node that generates the log entry. The software appends a colon (:) and a space to the string and it is inserted in the syslog message after the date stamp and before the syslog message content. Only one string can be entered. If multiple strings are entered, the last string overwrites the previous string. The alphanumeric string can contain lowercase (a-z), uppercase (A-Z) and numeric (0-9) characters. The no form of the command removes the log prefix string.
Default	log-prefix "TMNX".
Parameters	<i>log-prefix-string</i> — Specifies the alphanumeric string of up to 32 characters. Spaces and colons (:) cannot be used in the string.

level

Syntax	level <i>syslog-level</i> no level			
Context	config>service>vprn>log>syslog			
Description	This command configures the syslog message severity level threshold. All messages with severity level equal to or higher than the threshold are sent to the syslog target host. Only a single threshold level can be specified. If multiple levels are entered, the last level entered will overwrite the previously entered commands. The no form of the command reverts to the default value.			
Default	level info			
Parameters	<i>syslog-</i> — The threshold severity level name.			
	Values	emergency, alert, critical, error, warning, notice, info, debug		
	Router severity level	Numerical Severity (highest to lowest)	Configured Severity	Definition
		0	emergency	system is unusable

Router severity level	Numerical Severity (highest to lowest)	Configured Severity	Definition
3	1	alert	action must be taken immediately
4	2	critical	critical condition
5	3	error	error condition
6	4	warning	warning condition
	5	notice	normal but significant condition
1 cleared 2 indeterminate	6	info	informational messages
	7	debug	debug-level messages

port

Syntax	port <i>value</i> no port
Context	config>service>vprn>log>syslog
Description	This command configures the UDP port that will be used to send syslog messages to the syslog target host. The port configuration is needed if the syslog target host uses a port other than the standard UDP syslog port 514. Only one port can be configured. If multiple port commands are entered, the last entered port overwrites the previously entered ports. The no form of the command reverts to default value.
Default	no port
Parameters	<i>value</i> — The value is the configured UDP port number used when sending syslog messages. Values 1 to 65535

snmp-trap-group

Syntax	[no] snmp-trap-group <i>log-id</i>
Context	config>service>vprn>log
Description	This command creates the context to configure a group of SNMP trap receivers and their operational parameters for a given log-id.

A group specifies the types of SNMP traps and specifies the log ID which will receive the group of SNMP traps. A trap group must be configured in order for SNMP traps to be sent.

To suppress the generation of all alarms and traps see the **event-control** command. To suppress alarms and traps that are sent to this log-id, see the **filter** command. Once alarms and traps are generated they can be directed to one or more SNMP trap groups. Logger events that can be forwarded as SNMP traps are always defined on the main event source.

The **no** form of the command deletes the SNMP trap group.

Default	There are no default SNMP trap groups.
Parameters	<i>log-id</i> — The log ID value of a log configured in the log-id context. Alarms and traps cannot be sent to the trap receivers until a valid <i>log-id</i> exists.
Values	1 to 99

trap-target

Syntax	trap-target <i>name</i> [address <i>ip-address</i>] [port <i>port</i>] [snmpv1 snmpv2c snmpv3] notify-community <i>communityName</i> <i>snmpv3SecurityName</i> [security-level { no-auth-no-privacy auth-no-privacy privacy }] [replay] no trap-target <i>name</i>
Context	config>service>vprn>log>snmp-trap-group
Description	This command adds/modifies a trap receiver and configures the operational parameters for the trap receiver. A trap reports significant events that occur on a network device such as errors or failures.

Before an SNMP trap can be issued to a trap receiver, the [log-id](#), [snmp-trap-group](#) and at least one [snmp-trap-group](#) must be configured.

The [snmp-trap-group](#) command is used to add/remove a trap receiver from an [snmp-trap-group](#). The operational parameters specified in the command include:

- The IP address of the trap receiver
- The UDP port used to send the SNMP trap
- SNMP version
- SNMP community name for SNMPv1 and SNMPv2c receivers.
- Security name and level for SNMPv3 trap receivers.

A single **snmp-trap-group** *log-id* can have multiple trap-receivers. Each trap receiver can have different operational parameters.

An address can be configured as a trap receiver more than once as long as a different port is used for each instance.

To prevent resource limitations, only configure a maximum of 10 trap receivers.

If the same **trap-target name port port** parameter value is specified in more than one SNMP trap group, each trap destination should be configured with a different *notify-community* value. This allows a trap receiving an application, such as NMS, to reconcile a separate event sequence number stream for each router event log when multiple event logs are directed to the same IP address and port destination.

The **no** form of the command removes the SNMP trap receiver from the SNMP trap group.

Default No SNMP trap targets are defined.

Parameters *name* — specifies the name of the trap target up to 28 characters in length
address ip-address — The IP address of the trap receiver in dotted decimal notation.
Only one IP address destination can be specified per trap destination group.

Values

ipv4-address	a.b.c.d (host bits must be 0)
ipv6-address	x:x:x:x:x:x[-interface]
	x:x:x:x:x:x.d.d.d[-interface]
	x: [0 to FFFF]H
	d: [0 to 255]D
	interface: 32 characters maximum, mandatory for link local addresses

The ipv6-address applies to the 7750 SR.

port port — The destination UDP port used for sending traps to the destination, expressed as a decimal integer. Only one port can be specified per **trap-target** statement. If multiple traps need to be issued to the same address then multiple ports must be configured.

Values 1 to 65535

Default 162

snmpv1 | snmpv2c | snmpv3 — Specifies the SNMP version format to use for traps sent to the trap receiver.

The keyword **snmpv1** selects the SNMP version 1 format. When specifying **snmpv1**, the **notify-community** must be configured for the proper SNMP community string that the trap receiver expects to be present in alarms and traps messages. If the SNMP version is changed from **snmpv3** to **snmpv1**, then the **notify-community** parameter must be changed to reflect the community string rather than the *security-name* that is used by **snmpv3**.

The keyword **snmpv2c** selects the SNMP version 2c format. When specifying **snmpv2c**, the **notify-community** must be configured for the proper SNMP community string that the trap receiver expects to be present in alarms and traps messages. If the SNMP version is changed from **snmpv3** to **snmpv2c**, then the **notify-community** parameter must be changed to reflect the community string rather than the *security-name* that is used by **snmpv3**.

The keyword **snmpv3** selects the SNMP version 3 format. When specifying **snmpv3**, the **notify-community** must be configured for the SNMP *security-name*. If the SNMP version is changed from **snmpv1** or **snmpv2c** to **snmpv3**, then the **notify-community** parameter must be changed to reflect the *security-name* rather than the community string used by **snmpv1** or **snmpv2c**.

Pre-existing conditions are checked before the `snmpv3SecurityName` is accepted. These are:

- The user name must be configured.
- The v3 access group must be configured.
- The v3 notification view must be configured.

Values `snmpv1`, `snmpv2c`, `snmpv3`

Default `snmpv3`

notify-community community | security-name — Specifies the community string for **snmpv1** or **snmpv2c** or the **snmpv3** *security-name*. If no **notify-community** is configured, then no alarms nor traps will be issued for the trap destination. If the SNMP version is modified, the **notify-community** must be changed to the proper form for the SNMP version.

community — The community string as required by the **snmpv1** or **snmpv2c** trap receiver. The community string can be an ASCII string up to 31 characters in length.

security-name — The *security-name* as defined in the `config>system>security>user` context for SNMP v3. The *security-name* can be an ASCII string up to 31 characters in length.

security-level {no-auth-no-privacy | auth-no-privacy | privacy} — Specifies the required authentication and privacy levels required to access the views configured on this node when configuring an **snmpv3** trap receiver.

The keyword **no-auth-no-privacy** specifies no authentication and no privacy (encryption) are required.

The keyword **auth-no-privacy** specifies authentication is required but no privacy (encryption) is required. When this option is configured the *security-name* must be configured for **authentication**.

The keyword **privacy** specifies both authentication and privacy (encryption) is required. When this option is configured the *security-name* must be configured for **authentication** and **privacy**.

Values `no-auth-no-privacy`, `auth-no-privacy`, `privacy`

Default `no-auth-no-privacy`. This parameter can only be configured if SNMPv3 is also configured.

replay — Enable replay of missed events to target. If replay is applied to an SNMP trap target address, the address is monitored for reachability. Reachability is determined by whether or not there is a route in the routing table by which the target address can be reached. Before sending a trap to a target address, the SNMP module asks the PIP module if there is either an in-band or out-of-band route to the target address. If there is no route to the SNMP target address, the SNMP module saves the sequence-id of the first event that will be missed by the trap target. When the routing

table changes again so that there is now a route by which the SNMP target address can be reached, the SNMP module replays (for example, retransmits) all events generated to the SNMP notification log while the target address was removed from the route table. Because of route table change convergence time, it is possible that one or more events may be lost at the beginning or end of a replay sequence. The cold-start-wait and route-recovery-wait timers under config>log>app-route-notifications can help reduce the probability of lost events.

3.8.2.8 Remote Server Commands

aaa

Syntax	aaa
Context	config>service>vprn
Description	This command enters the context to configure AAA on the VPRN.

remote-servers

Syntax	remote-servers
Context	config>service>vprn>aaa
Description	This command enters the context to configure AAA remote servers on the VPRN.

radius

Syntax	radius [create] no radius
Context	config>service>vprn>aaa>remote-servers
Description	This command creates the context to configure RADIUS authentication on the VPRN. Implement redundancy by configuring multiple server addresses for each VPRN. The no form of the command removes the RADIUS configuration.

access-algorithm

Syntax	access-algorithm {direct round-robin}

no access-algorithm

Context	config>service>vprn>aaa>remote-servers>radius
Description	This command indicates the algorithm used to access the set of RADIUS servers.
Default	access-algorithm direct
Parameters	direct — The first server will be used as primary server for all requests, the second as secondary and so on. round-robin — The first server will be used as primary server for the first request, the second server as primary for the second request, and so on. If the router gets to the end of the list, it starts again with the first server.

accounting

Syntax	[no] accounting
Context	config>service>vprn>aaa>remote-servers>radius
Description	This command enables RADIUS accounting. The no form of this command disables RADIUS accounting.
Default	no accounting

accounting-port

Syntax	accounting-port <i>port</i> no accounting-port
Context	config>service>vprn>aaa>remote-servers>radius
Description	This command specifies a UDP port number on which to contact the RADIUS server for accounting requests.
Default	accounting-port 1813
Parameters	<i>port</i> — Specifies the UDP port number. Values 1 to 65535 Default 1813

authorization

Syntax	[no] authorization

Context	config>service>vprn>aaa>remote-servers>radius
Description	This command configures RADIUS authorization parameters for the system.
Default	no authorization

interactive-authentication

Syntax	[no] interactive-authentication
Context	config>service>vprn>aaa>remote-servers>radius
Description	This command enables RADIUS interactive authentication for the system. Enabling interactive-authentication forces RADIUS to fall into challenge/response mode.
Default	no interactive-authentication

port

Syntax	port <i>port</i> no port
Context	config>service>vprn>aaa>remote-servers>radius
Description	This command configures the TCP port number to contact the RADIUS server. The no form of the command reverts to the default value.
Default	port 1812 (as specified in RFC 2865, <i>Remote Authentication Dial In User Service (RADIUS)</i>)
Parameters	<i>port</i> — Specifies the TCP port number to contact the RADIUS server. Values 1 to 65535

retry

Syntax	retry <i>count</i> no retry
Context	config>service>vprn>aaa>remote-servers>radius
Description	This command configures the number of times the router attempts to contact the RADIUS server for authentication if there are problems communicating with the server. The no form of the command reverts to the default value.
Default	retry 3

Parameters *count* — Specifies the retry count.

Values 1 to 10

server

Syntax **server** *index* **address** *ip-address* **secret** *key* [*hash* | *hash2*]
no server *index*

Context config>service>vprn>aaa>remote-servers>radius

Description This command adds a RADIUS server and configures the RADIUS server IP address, index, and key values.

Up to five RADIUS servers can be configured at any one time. RADIUS servers are accessed in order from lowest to highest index for authentication requests until a response from a server is received. A higher indexed server is only queried if no response is received from a lower indexed server (which implies that the server is not available). If a response from a server is received, no other RADIUS servers are queried. It is assumed that there are multiple identical servers configured as backups and that the servers do not have redundant data.

The **no** form of the command removes the server from the configuration.

Default no server

Parameters *index* — Specifies the index for the RADIUS server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index.

Values 1 to 5

ip-address — Specifies the IP address of the RADIUS server. Two RADIUS servers cannot have the same IP address. An error message is generated if the server address is a duplicate.

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..FFFF]H
	d: [0..255]D

key — Specifies the secret key to access the RADIUS server. This secret key must match the password on the RADIUS server.

Values Up to 64 characters in length.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

timeout

Syntax	timeout <i>seconds</i> no timeout
Context	config>service>vprn>aaa>remote-servers>radius
Description	This command configures the number of seconds the router waits for a response from a RADIUS server. The no form of the command reverts to the default value.
Default	timeout 3
Parameters	<i>seconds</i> — Specifies the number of seconds the router waits for a response from a RADIUS server, expressed as a decimal integer. Values 1 to 90

use-default-template

Syntax	[no] use-default-template
Context	config>service>vprn>aaa>remote-servers>radius
Description	This command specifies whether the RADIUS default user template is actively applied to the RADIUS user if no VSAs are returned with the auth-accept from the RADIUS server. When enabled, the radius_default user-template is actively applied if no VSAs are returned with the auth-accept from the RADIUS server and radius authorization is enabled. The no form of the command disables the use of the RADIUS default template.
Default	no use-default-template

tacplus

Syntax	no tacplus tacplus create
Context	config>service>vprn>aaa>remote-servers
Description	This command creates the context to configure TACACS+ authentication on the VPRN. Configure multiple server addresses for each router for redundancy. The no form of the command removes the TACACS+ configuration.

accounting

Syntax	accounting [record-type {start-stop stop-only}] no accounting
Context	config>service>vprn>aaa>remote-servers>tacplus
Description	This command configures the type of accounting record packet that is to be sent to the TACACS+ server. The record-type parameter indicates whether TACACS+ accounting start and stop packets be sent or just stop packets be sent.
Default	no accounting
Parameters	record-type start-stop — Specifies that a TACACS+ start packet is sent whenever the user executes a command and a TACACS+ stop packet when command execution is complete. record-type stop-only — Specifies that only a TACACS+ stop packet is sent whenever the command execution is complete.

authorization

Syntax	authorization [use-priv-lvl] no authorization
Context	config>service>vprn>aaa>remote-servers>tacplus
Description	This command configures TACACS+ authorization parameters for the VPRN.
Default	no authorization

Parameters *use-priv-lvl* — Automatically performs a single authorization request to the TACACS+ server for *cmd** (all commands) immediately after login, and then use the local profile associated (via the *priv-lvl-map*) with the *priv-lvl* returned by the TACACS+ server for all subsequent authorization (except *enable-admin*). After the initial authorization for *cmd**, no further authorization requests will be sent to the TACACS+ server (except *enable-admin*).

interactive-authentication

Syntax **[no] interactive-authentication**

Context config>service>vprn>aaa>remote-servers>tacplus

Description This configuration instructs the SR OS to send no username nor password in the TACACS+ start message, and to display the *server_msg* in the GETUSER and GETPASS response from the TACACS+ server. Interactive authentication can be used to support a One Time Password scheme (e.g. S/Key). An example flow (e.g. with a telnet connection) is as follows:

- The SR OS will send an authentication start request to the TACACS+ server with no username nor password.
- TACACS+ server replies with TAC_PLUS_AUTHEN_STATUS_GETUSER and a *server_msg*.
- The SR OS displays the *server_msg*, and collects the user name.
- The SR OS sends a continue message with the user name.
- TACACS+ server replies with TAC_PLUS_AUTHEN_STATUS_GETPASS and a *server_msg*.
- The SR OS displays the *server_msg* (which may contain, for example, an S/Key for One Time Password operation), and collects the password.
- The SR OS sends a continue message with the password.
- TACACS+ server replies with PASS or FAIL.

When interactive-authentication is disabled the SR OS will send the username and password in the *tacplus* start message. An example flow (e.g. with a telnet connection) is as follows:

- TAC_PLUS_AUTHEN_TYPE_ASCII.
 - the login username in the “user” field.
 - the password in the *user_msg* field (while this is non-standard, it does not cause interoperability problems).
- TACACS+ server ignores the password and replies with TAC_PLUS_AUTHEN_STATUS_GETPASS.
- The SR OS sends a continue packet with the password in the *user_msg* field.
- TACACS+ server replies with PASS or FAIL.

When interactive-authentication is enabled, *tacplus* must be the first method specified in the authentication-order configuration.

Default no interactive-authentication

priv-lvl-map

Syntax [no] priv-lvl-map

Context config>service>vprn>aaa>remote-servers>tacplus

Description This command enables the context to specify a series of mappings between TACACS+ priv-lvl and locally configured profiles for authorization. These mappings are used when the use-priv-lvl option is specified for tacplus authorization.

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

Default priv-lvl-map

priv-lvl

Syntax priv-lvl priv-lvl user-profile-name
no priv-lvl priv-lvl

Context config>service>vprn>aaa>remote-servers>tacplus>priv-lvl-map

Description This command maps a specific TACACS+ priv-lvl to a locally configured profile for authorization. This mapping is used when the **use-priv-lvl** option is specified for TACPLUS authorization.

Parameters *priv-lvl* — Specifies the privilege level used when sending a TACACS+ ENABLE request.

Values 0 to 15

user-profile-name — Specifies the user profile for this mapping.

server

Syntax server index address ip-address secret key [{hash | hash2}][port port]
no server index

Context config>service>vprn>aaa>remote-servers>tacplus

Description This command adds a TACACS+ server and configures the TACACS+ server IP address, index, and key values.

Up to five TACACS+ servers can be configured at any one time. TACACS+ servers are accessed in order from lowest index to the highest index for authentication requests.

The **no** form of the command removes the server from the configuration.

Default	No TACACS+ servers are configured.										
Parameters	<i>index</i> — Specifies the index for the TACACS+ server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from the lowest index to the highest index. Values 1 to 5 <i>ip-address</i> — Specifies the IP address of the TACACS+ server. Two TACACS+ servers cannot have the same IP address. An error message is generated if the server address is a duplicate. Values <table> <tr> <td>ipv4-address</td><td>a.b.c.d (host bits must be 0)</td></tr> <tr> <td>ipv6-address</td><td>x:x:x:x:x:x:x (eight 16-bit pieces)</td></tr> <tr> <td></td><td>x:x:x:x:x:d.d.d.d</td></tr> <tr> <td></td><td>x: [0..FFFF]H</td></tr> <tr> <td></td><td>d: [0..255]D</td></tr> </table> secret key — Specifies the secret key to access the RADIUS server. This secret key must match the password on the RADIUS server. Values Up to 128 characters in length. hash — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the hash2 encrypted variable cannot be copied and pasted. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified. <i>port</i> — Specifies the port ID. Values 0 to 65535	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..FFFF]H		d: [0..255]D
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..FFFF]H										
	d: [0..255]D										

shutdown

Syntax	[no] shutdown
Context	config>service>vprn>aaa>remote-servers>tacplus
Description	This command administratively disables the TACACS+ protocol operation. Shutting down the protocol does not remove or change the configuration other than the administrative state. 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 the command administratively enables the protocol which is the default state.

Default no shutdown

timeout

Syntax **timeout** *seconds*
no timeout

Context config>service>vprn>aaa>remote-servers>tacplus

Description This command configures the number of seconds the router waits for a response from a TACACS+ server.

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

Default timeout 3

Parameters *seconds* — Specifies the number of seconds the router waits for a response from a TACACS+ server, expressed as a decimal integer.

Values 1 to 90

use-default-template

Syntax [**no**] **use-default-template**

Context config>system>security>tacplus

Description This command specifies whether the tacplus_default user-template is actively applied to the TACACS+ user. When enabled, the tacplus_default user-template is actively applied if tacplus authorization is enabled (without the use-priv-lvl option).

The **no** form of the command disables the feature.

Default use-default-template

3.8.2.9 VPRN Management Commands

management

Syntax **management** [**create**]
no management

Context config>service>vprn

Description	This command enters the node management configuration within VPRN.
Parameters	create — Creates a management server entry.

allow-ftp

Syntax	[no] allow-ftp
Context	config>service>vprn>management
Description	This commands allows access to the FTP server from VPRN. The no form of this command removes FTP access for this VPRN.

allow-ssh

Syntax	[no] allow-ssh
Context	config>service>vprn>management
Description	This command allows configuration of the SSH parameters. The no form of this command disallows configuration of the SSH parameters.

allow-telnet

Syntax	[no] allow-telnet
Context	config>service>vprn>management
Description	This command allows access to the Telnet server from a VPRN. The no form of this command removes the Telnet access.

allow-telnet6

Syntax	[no] allow-telnet6
Context	config>service>vprn>management
Description	This command allows access to the Telnet IPv6 server from a VPRN. The no form of this command removes the Telnet IPv6 access.

3.8.2.10 Multicast VPN Commands

mvpn

Syntax	mvpn
Context	config>service>vprn
Description	This command enters the context to configure MVPN-related parameters for the IP VPN.

auto-discovery

Syntax	auto-discovery [default mdt-safi] [source-address <i>ip-address</i>]
Context	config>service>vprn>mvpn
Description	This command enables MVPN membership auto-discovery through BGP. When auto-discovery is enabled, PIM peering on the inclusive provider tunnel is disabled. Changing auto-discovery configuration requires shutdown of this VPRN instance. The no form of the command disables MVPN membership auto-discovery through BGP.
Default	default
Parameters	default — Enables AD route exchange based on format defined in NG-MVPN (RFC6514). mdt-safi — Enables AD route exchange based on mdt-safi format defined in <i>draft-rosen-vpn-mcast</i>. This command optionally specifies a source-address - an IP address to be used by Rosen MVPN or NG-MVPN for core diversity, non-default IGP instances (not using system IP). Two unique IP addresses for PIM or GRE MVPNs are supported. The two unique IP address restriction does not apply to MVPNs with MPLS tunnels (for example, RSVP and MLDP). For instances using default System IP, source address configuration should not be specified to avoid consuming one of the addresses. Explicitly defining a source-address allows GRE-encapsulated Rosen MVPN or NG-MVPN multicast traffic (Default and Data MDT) to originate from a configured IP address, so the source IP address of the GRE packets will not be the default system IP address. Value: <i>ip-address</i> — An IPv4 address. To achieve the desired functionality the address should be a pre-configured non-default ISIS or OSPF loopback address for an IGP instance using loopback address different from the system IP loopback.

c-mcast-signaling

Syntax	c-mcast-signaling {bgp pim} no c-mcast-signaling
Context	config>service>vprn>mvpn
Description	This command specifies BGP or PIM, for PE-to-PE signaling of CE multicast states. When this command is set to PIM and neighbor discovery by BGP is disabled, PIM peering will be enabled on the inclusive tree. Changes may only be made to this command when the mvpn node is shutdown. The no form of the command reverts it back to the default.
Default	c-mcast-signaling bgp
Parameters	bgp — Specifies to use BGP for PE-to-PE signaling of CE multicast states. Auto-discovery must be enabled. pim — Specifies to use PIM for PE-to-PE signaling of CE multicast states.

intersite-shared

Syntax	intersite-shared [persistent-type5-adv] [kat-type5-adv-withdraw] no intersite-shared
Context	config>service>vprn>mvpn
Description	This command specifies whether to use inter-site shared C-trees or not. Optional parameters allow enabling additional inter-site shared functionality. Not specifying an optional parameter when executing the command disables that parameter.
Default	n/a
Parameters	persistent-type5-adv — When specified for inter-site shared trees enabled, this parameter ensures that Type 5 SA routes are generated for the multicast source even if no joins are present for that source. When the parameter is not specified, the Type 5 SA routes are withdrawn where the prune from the last receiver is received for the multicast source. kat-type5-adv-withdraw — When specified for inter-site shared trees, this parameter allows operators to enable KeepAlive Timers (KAT) on source PEs for ng-MVPN inter-site shared deployments. On a multicast source failure, a KAT expiry on source PEs will trigger a withdrawal of Type-5 Source-Active (S-A) route and switch from (C-S,C-G) to (C-*,C-G). When receiver PEs process reflected Type-5 S-A route withdrawals, they will withdraw their Type-7 ng-MVPN routes to the failed multicast source. The following conditions apply: • KAT must only be enabled on source PEs. • Functionality is supported with mLDP and RSVP-TE in the P-instance.

- Local receiver per (C-S, C-G) must be configured on source PEs running KAT.

mdt-type

Syntax	mdt-type { sender-receiver sender-only receiver-only }
Context	config>service>vprn>mvpn
Description	This command allows restricting MVPN instance per PE node to a specific role. By default, MVPN instance on a given PE node assumes the role of being a sender as well as receiver. This creates a mesh of MDT/PMSI across all PE nodes from this PE. This command provides an option to configure either a sender-only or receiver-only mode per PE node. Restricting the role of a PE node prevents creating full mesh of MDT/PMSI across all PE nodes that are participating in MVPN instance. auto-rp-discovery cannot be enabled together with mdt-type sender-only or mdt-type receiver-only, or wildcard-spmsi configurations. The no version of this command restores the default (sender-receiver).
Default	mdt-type sender-receiver
Parameters	sender-receiver — MVPN has both sender and receivers connected to PE node. sender-only — MVPN has only senders connected to PE node. receiver-only — MVPN has only receivers connected to PE node.

red-source-list

Syntax	red-source-list
Context	config>service>vprn>mvpn
Description	This command enables context to configure list of redundant source prefixes for preferred source selection.

src-prefix

Syntax	src-prefix <i>ip-address/mask</i> [<i>ip-address/mask</i>] no src-prefix <i>ip-address/mask</i>
Context	config>service>vprn>mvpn>red-source-list
Description	This command configures multicast source IPv4 prefixes for preferred source selection. Single or multi-line inputs are allowed.

The **no** form of the command deletes specified prefix from the list.

Default	No prefixes are specified.
Parameters	<i>ip-address/mask</i> — IPv4 address prefix with mask. Up to 8 maximum.

ipv6

Syntax	ipv6
Context	config>service>vprn>mvpn>red-source-list
Description	This command enables context to configure list of redundant IPv6 source prefixes for preferred source selection.

src-prefix

Syntax	src-prefix <i>ipv6-ip-address/prefix-length</i> [<i>ipv6-address/prefix-length</i>] no <i>ipv6-ip-address/prefix-length</i>
Context	config>service>vprn>mvpn>red-source-list>ipv6
Description	This command configures multicast source IPv6 prefixes for preferred source selection. Single or multi-line inputs are allowed. The no form of the command deletes specified prefix from the list
Default	No prefixes are specified.
Parameters	<i>ipv6-ip-address/mask</i> — IPv6 address prefix with prefix-length. Up to 8 maximum.

rpf-select

Syntax	rpf-select
Context	config>service>vprn>mvpn
Description	This command enables context for VRF extranet mapping for C-instance receivers in this receiver MVPN instance to multicast streams in P-instance core MVPN instances.

core-mvpn

Syntax	[no] core-mvpn <i>service-id</i>
Context	config>service>vprn>mvpn>rpf-select

Description	This command enables context for VRF extranet mapping for C-instance receivers in this receiver MVPN instance to multicast streams in the specified P-instance core MVPN instance.

group-prefix

Syntax	group-prefix <i>ip-address/mask</i> [<i>ip-address/mask</i>] [starg] no group-prefix <i>ip-address/mask</i>
Context	config>service>vprn>mvpn>rpf-select>core-mvpn
Description	This command configures multicast group IPv4 prefixes for the MVPN with per-group mapping extranet functionality. Multiple lines are allowed. Duplicate prefixes are ignored. When the starg option is specified, extranet functionality is enabled for PIM ASM as for the specified group. When the option is not specified (not recommended with PIM ASM), the PIM ASM join will be mapped and data plane will be established, but the control plane will not be updated on SPT switchover, unless the switchover is driven by a CPE router on a receiver side. The no form of the command deletes specified prefix from the list, or removes mapping of all prefixes if group-prefix any was specified.
Parameters	<i>ip-address/mask</i> — Specifies the IPv4 multicast address prefix with mask. Up to 8 addresses can be specified in a single statement.

provider-tunnel

Syntax	provider-tunnel
Context	config>service>vprn>mvpn
Description	This command enables context to configure tunnel parameters for the MVPN.

inclusive

Syntax	inclusive
Context	config>service>vprn>mvpn>provider-tunnel
Description	This command enters the context for specifying inclusive provider tunnels.

bier

Syntax	[no] bier

Context	config>service>vprn>mvpn>provider-tunnel>inclusive config>service>vprn>mvpn>provider-tunnel>selective
Description	This command creates a BIER inclusive or selective provider tunnel. The no form of this command deletes the tunnel.

shutdown

Syntax	shutdown no shutdown
Context	config>service>vprn>mvpn>provider-tunnel>inclusive>bier config>service>vprn>mvpn>provider-tunnel>selective>bier
Description	This command administratively disables and enables use of BIER for the provider tunnel.
Default	no shutdown

sub-domain

Syntax	sub-domain <i>sub-domain</i> no sub-domain
Context	config>service>vprn>mvpn>provider-tunnel>inclusive>bier config>service>vprn>mvpn>provider-tunnel>selective>bier
Description	This command sets the sub-domain used to attach the BIER provider tunnel. Both PMSI within the MVPN need to have the same sub-domain. The no form of this command removes the sub-domain.
Parameters	<i>sub-domain</i> — The identifier of the sub-domain. Values 0 to 255

bsr

Syntax	bsr {unicast spmsi} no bsr
Context	config>service>vprn>mvpn>pt>inclusive
Description	This command configures the type of BSR signaling used. The no form of the command restores the default.
Default	no bsr

Parameters	unicast — BSR PDUs are sent/forwarded using unicast PDUs (default). spmsi — BSR PDUs are sent/forwarded using S-PMSI full mesh.

mldp

Syntax	mldp no mldp
Context	config>service>vprn>mvpn>pt>inclusive
Description	This command enables use of mLDP LSP for the provider tunnel.
Default	no mldp

shutdown

Syntax	shutdown no shutdown
Context	config>service>vprn>mvpn>pt>inclusive>mldp
Description	This command administratively disables and enables use of mLDP LSP for the provider tunnel.
Default	no shutdown

pim

Syntax	pim {asm ssm} grp-ip-address no pim
Context	config>service>vprn>mvpn>pt>inclusive
Description	This command specifies the PIM mode to use, ASM or SSM, for PIM-based inclusive provider tunnels and the multicast group address to use. Also enables the context for specifying parameters for PIM peering on the inclusive provider tunnel. Auto-discovery must be enabled in order for SSM to operate. The no form of the command removes the pim context including the statements under the context.
Default	no pim
Parameters	asm — Specifies to use PIM ASM for inclusive provider tunnels. ssm — Specifies to u PIM SSM for inclusive provider tunnels.

group-address — Specifies the multicast group address to use.

hello-interval

Syntax	hello-interval <i>hello-interval</i> no hello-interval
Context	config>service>vprn>mvpn>pt>inclusive>pim
Description	This command configures the frequency at which PIM Hello messages are transmitted on this interface. The no form of this command resets the configuration to the default value.
Default	hello-interval 30
Parameters	<i>hello-interval</i> — Specifies the hello interval in seconds. A 0 (zero) value disables the sending of Hello messages (the PIM neighbor will never timeout the adjacency). Values 0 to 255 seconds

hello-multiplier

Syntax	hello-multiplier <i>deci-units</i> no hello-multiplier
Context	config>service>vprn>mvpn>pt>inclusive>pim
Description	This command configures the multiplier to determine the holdtime for a PIM neighbor on this interface. The hello-multiplier in conjunction with the hello-interval determines the holdtime for a PIM neighbor.
Parameters	<i>deci-units</i> — Specify the value, specified in multiples of 0.1, for the formula used to calculate the holdtime based on the hello-multiplier : (hello-interval * hello-multiplier) / 10 This allows the PIMv2 default hello-multiplier of 3.5 and the default timeout of 105 seconds to be supported. Values 20 to 100 Default 35

improved-assert

Syntax	[no] improved-assert

Context	config>service>vprn>mvpn>pt>inclusive>pim
Description	This command enables improved assert procedure on the PIM inclusive provider tunnel. The no form of the command disables improved assert procedure.
Default	enabled

three-way-hello

Syntax	[no] three-way-hello
Context	config>service>vprn>mvpn>pt>inclusive>pim
Description	This command enables PIM three-way hello on the inclusive provider tunnel. The no form of the command disables the PIM three-way hello.
Default	disabled

tracking-support

Syntax	[no] tracking-support
Context	config>service>vprn>mvpn>pt>inclusive>pim
Description	This command enables the setting of the T bit in the LAN Prune Delay option of the Hello message. This indicates the router's capability to disable Join message suppression. The no form of the command disables the setting.
Default	disabled

rsvp

Syntax	rsvp no rsvp
Context	config>service>vprn>mvpn>pt>inclusive
Description	This command enters the context for specifying RSVP P2MP LSP for the provider tunnel. The no form of the command removes the rsvp context including all the statements in the context.
Default	no rsvp

enable-bfd-root

Syntax	enable-bfd-root [transmit-interval] [multiplier <i>multiplier</i>] no enable-bfd-root
Context	config>service>vprn>mvpn>pt>inclusive>rsvp
Description	This command enables unidirectional multi-point BFD session on a sender (Root) PE node for upstream fast failure detection over RSVP-TE P2MP LSP.
Parameters	transmit-interval — Sets the transmit interval, in milliseconds. Values 10 to 100000 Default 100 multiplier <i>multiplier</i> — Sets the multiplier for the BFD session. Values 3 to 20 Default 3

enable-bfd-leaf

Syntax	[no] enable-bfd-leaf
Context	config>service>vprn>mvpn>pt>inclusive>rsvp
Description	This command enables unidirectional multi-point BFD session on a receiver (leaf) PE node for upstream fast failure detection over RSVP-TE P2MP LSP.

lsp-template

Syntax	lsp-template no lsp-template
Context	Context config>service>vprn>mvpn>pt>inclusive>rsvp
Description	This command specifies the use of automatically created P2MP LSP as the provider tunnel. The P2MP LSP will be signaled using the parameters specified in the template, such as bandwidth constraints, and so on.
Default	no lsp-template

shutdown

Syntax	shutdown no shutdown

Context	config>service>vprn>mvpn>pt>inclusive>rsvp>lsp-template
Description	This command administratively disables and enables use of RSVP P2MP LSP for the provider tunnel.
Default	no shutdown

wildcard-spmsi

Syntax	wildcard-spmsi no wildcard-spmsi
Context	config>service>vprn>mvpn>pt>inclusive
Description	This command enables RFC 6625 (C-*, C-*) S-PMSI functionality for NG-MVPN. When enabled, (C-*, C-*) S-PMSI is used instead of I-PMSI for this MVPN. Wildcard S-PMSI uses the I-PMSI LSP template. auto-rp-discovery cannot be enabled together with mdt-type sender-only or mdt-type receiver-only, or wildcard-spmsi configurations. The no form disables the (C-*, C-*) S-PMSI functionality.
Default	no wildcard-spmsi

selective

Syntax	selective
Context	config>service>vprn>mvpn>provider-tunnel
Description	This command enters the context to specify selective provider tunnel parameters.

auto-discovery-disable

Syntax	[no] auto-discovery-disable
Context	config>service>vprn>mvpn>pt>selective
Description	This command disables C-trees to P-tunnel binding auto-discovery through BGP so it is signaled using PIM join TLVs. This command requires the c-mcast-signaling parameter to be set to PIM. For multi-stream S-PMSI, this command must be enabled for BGP auto-discovery to function. The no form of the command enables multicast VPN membership auto-discovery through BGP.

Default auto-discovery-disable

data-delay-interval

Syntax **data-delay-interval** *value*
no data-delay-interval

Context config>service>vprn>mvpn>pt>selective

Description This command specifies the interval, in seconds, before a PE router connected to the source switches traffic from the inclusive provider tunnel to the selective provider tunnel.

This command is not applicable to multi-stream S-PMSI,

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

Default 3 seconds

Parameters *value* — Specifies the data delay interval, in seconds.

Values 3 to 180

data-threshold

Syntax **data-threshold** {*c-grp-ip-addr/mask* | *c-grp-ip-addr netmask*} *s-pmsi-threshold*
[**pe-threshold-add** *pe-threshold-add*] [**pe-threshold-delete** *pe-threshold-delete*]
data-threshold *c-grp-ipv6-addr/prefix-length* *s-pmsi-threshold* [**pe-threshold-add** *pe-threshold-add*]
[**pe-threshold-delete** *pe-threshold-delete*]
no data-threshold {*c-grp-ip-addr/mask* | *c-grp-ip-addr netmask*}
no data-threshold *c-grp-ipv6-addr/prefix-length*

Context config>service>vprn>mvpn>pt>selective

Description This command specifies the data rate threshold that triggers the switch from the inclusive provider tunnel to the selective provider tunnel for (C-S, C-G) within the group range. Optionally, PE thresholds for creating/deleting ng-MVPN S-PMSI may also be specified. Omitting the PE thresholds, preserves currently set value (or defaults if never set). Multiple statements (one per a unique group) are allowed in the configuration.

This command is not applicable to multi-stream S-PMSI,

The **no** form of the command removes the values from the configuration.

Default no data-threshold

Parameters *group-address/mask* — Specifies a multicast group address and netmask length.
c-grp-ip-addr/mask | *c-grp-ip-addr netmask* — Specifies an IPv4 multicast group address and netmask length or network mask.

c-grp-ipv6-addr/prefix-length — Specifies an IPv6 multicast group address and prefix length.

s-pmsi-threshold — Specifies the rate, in kb/s. If the rate for a (C-S, C-G)) within the specified group range exceeds the threshold, traffic for the (C-S, C-G) will be switched to the selective provider tunnel.

s-pmsi-threshold-add — Specifies the number of receiver PEs for creating S-PMSI. When the number of receiver PEs for a given multicast group configuration is non-zero and below the threshold and BW threshold is satisfied, S-PMSI is created.

s-pmsi-threshold-delete — Specifies the number of receiver PEs for deleting S-PMSI. When the number of receiver PEs for a given multicast group configuration is above the threshold, S-PMSI is deleted and the multicast group is moved to I-PMSI or a wildcard S-PMSI. It is recommended that the delete threshold be significantly larger than the add threshold, to avoid re-signaling of S-PMSI as the receiver PE count fluctuates.

Values

c-grp-ip-addr : multicast group address a.b.c.d

mask [4 to 32]

netmask : a.b.c.d (network bits all 1 and host bits all 0)

s-pmsi-threshold : [1 to 4294967294](threshold in kb/s)

c-grp-ipv6-addr : multicast 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]

pe-threshold-add: [1 to 65535], if never specified, 65535 is used (add threshold always met)

pe-threshold-delete: [2 to 65535], if never specified, 65535 is used (delete threshold never met)

join-tlv-packing-disable

Syntax [no] join-tlv-packing-disable

Context config>service>vprn>mvpn>pt>selective

Description This command enables packing of MDT join TLVs into a single PDU to improve efficiency, if multiple join TLVs are available at the time of transmission.

The **no** form of the command disables packing of MDT join TLVs into a single PDU.

Default no join-tlv-packing-disable

multistream-spmsi

Syntax	[no] multistream-spmsi <i>index</i> [create]
Context	config>service>vprn>mvpn>pt>selective
Description	This command creates a multi-stream S-PMSI policy. Having multiple multi-stream S-PMSIs per MVPN creates a link list, in which the first match (lowest index) will be chosen for a multicast stream. The number of configured multi-stream S-PMSIs cannot exceed the configured maximum S-PMSI for a given MVPN.
Parameters	<i>index</i> — Specifies the index number.
Values	1 to 1024

group

Syntax	[no] group <i>ip-address</i> [/ <i>mask</i>]		
Context	config>service>vprn>mvpn>pt>selective>multistream-spmsi		
Description	This command creates group prefixes that map to the multicast stream. At least one source must be specified for the policy to be active.		
Parameters	<i>ip-address/mask</i> — Specifies the IP address.		
	Values		
	ipv4-prefix	a.b.c.d	
	ipv4-prefix-le		[0..32]
	ipv6-prefix		x:x:x:x:x:x:x (eight 16-bit pieces)
			x:x:x:x:x:d.d.d.d
		x	[0..FFFF]H
		d	[0..255]D
		ipv6-prefix-le	[0..128]

source

Syntax	[no] source <i>ip-address</i> [<i>/mask</i>] [no] source any
Context	config>service>vprn>mvpn>pt>selective>multistream-spmsi>group
Description	This command creates source prefixes for specific groups that map to the multicast stream.
Parameters	<i>ip-address/mask</i> — Specifies the IP address of the group.
Values	

ipv4-prefix	a.b.c.d	
	ipv4-prefix-le	[0..32]
ipv6-prefix		x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
	x	[0..FFFF]H
	d	[0..255]D
	ipv6-prefix-le	[0..128]

lsp-template

Syntax	[no] lsp-template <i>name</i>
Context	config>service>vprn>mvpn>pt>selective>multistream-spmsi
Description	This command creates a RSVP-TE LSP template for S-PMSI. Multi-stream S-PMSIs can share a single template or can each use their own template.
Parameters	<i>name</i> — Specifies the LSP template name, up to 32 characters.

mdt-pim

Syntax	mdt-pim mode {asm ssm} group-address <i>group-ip-address</i> no mdt-pim
Context	config>service>vprn>mvpn>pt>selective>multistream-spmsi
Description	This command creates a multi-stream MDT that could match many (C-S,C-G)s into a single data MDT.
Parameters	<i>group-ip-address</i> — Specifies the group address of this data MDT, that is the provider group address.

shutdown

Syntax	[no] shutdown
Context	config>service>vprn>mvpn>pt>selective>multistream-spmsi
Description	This commands enables multi-stream S-PSMI. At least one group must be active in a policy.

pim-asm

Syntax	[no] pim-asm { <i>grp-ip-address/mask</i> <i>grp-ip-address netmask</i> }

Context config>service>vprn>mvpn>pt>selective

Description This command specifies the range of PIM-ASM groups to use on the sender PE to setup ASM multicast tree for draft Rosen based Data MDT.

rsvp

Syntax [no] rsvp

Context config>service>vprn>mvpn>pt>inclusive
config>service>vprn>mvpn>pt>selective

Description This command enables use of P2MP RSVP as the inclusive or selective provider tunnel.

Default no rsvp

lsp-template

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

Context config>service>vprn>mvpn>pt>inclusive
config>service>vprn>mvpn>pt>selective>rsvp

Description This command specifies the use of automatically created P2MP LSP as the inclusive or selective provider tunnel. The P2MP LSP will be signaled using the parameters specified in the template, such as bandwidth constraints, and so on

Default no lsp-template

mldp

Syntax [no] mldp

Context config>service>vprn>mvpn>pt>inclusive
config>service>vprn>mvpn>pt>selective

Description This command enables use of P2MP mLDP LSP as inclusive or selective PMSI tunnels.

For multi-stream S-PMSI, either LDP or RSVP-TE must first be configured before multi-stream policy can be configured.

Default no mldp

maximum-p2mp-spmsi

Syntax [no] maximum-p2mp-spmsi

Context	config>service>vprn>mvpn>pt>selective
Description	This command specifies the maximum number of RSVP P2MP or LDP P2MP S-PMSI tunnels for the mVPN. When the limit is reached, no more RSVP P2MP S-PMSI or LDP P2MP S-PMSI is created and traffic over the data-threshold will stay on I-PMSI.
Default	10
Parameters	<i>number</i> — Specifies the maximum number of RSVP P2MP or LDP P2MP S-PMSI tunnel for the mVPN.
Values	1 to 4k
Default	10

shutdown

Syntax	[no] shutdown
Context	config>service>vprn>mvpn>pt>inclusive>rsvp>lsp-template config>service>vprn>mvpn>pt>inclusive>mldp config>service>vprn>mvpn>pt>selective>rsvp>lsp-template config>service>vprn>mvpn>pt>selective>mldp
Description	This command administratively disables/enables use of P2MP RSVP LSP template or mLDLP LSP for inclusive or selective PMSI tunnels.
Default	no shutdown

enable-asm-mdt

Syntax	[no] enable-asm-mdt
Context	config>service>vprn>mvpn>pt>selective
Description	This command enables Data MDT with PIM-ASM mode on the receiver PE node. PIM-ASM or PIM-SSM operation mode is derived based on the locally configured SSM range on the node. If asm-mode is disabled using this command, then PIM-SSM mode is enabled for all groups, independent of the configured SSM range on the node.

pim-ssm

Syntax	pim-ssm {grp-ip-address/mask grp-ip-address netmask} no pim-ssm
Context	config>service>vprn>mvpn>pt>selective

Description	This command specifies the PIM SSM groups to use for the selective provider tunnel.
Parameters	<i>group-address/mask</i> — Specifies a multicast group address and netmask length.

umh-pe-backup

Syntax	umh-pe-backup
Context	config>service>vprn>mvpn
Description	This command enables context to configure primary and standby upstream PE association for the MVPN.

umh-pe

Syntax	umh-pe <i>ip-address</i> standby <i>ip-address</i> no umh-pe <i>ip-address</i>
Context	config>service>vprn>mvpn>umh-pe-backup
Description	This command assigns a standby PE to each primary PE that must be selected as an alternative PE in case the UFD session on tunnel from primary PE is detected down. Standby for a PE cannot be modified without shutting down the MVPN instance. If a primary PE is not assigned a standby PE then the UMH selection would fall back to the default method.

umh-selection

Syntax	umh-selection { highest-ip hash-based tunnel-status unicast-rt-pref } no umh-selection
Context	config>service>vprn>mvpn
Description	This command specifies which UMH selection mechanism to use, highest IP address, hash based or provider tunnel status. The no form of the command resets it back to default.
Default	umh-selection highest-ip
Parameters	highest-ip — Specifies that the highest next-hop IP address is selected as UMH. The RTM may have just one next-hop to the source, but highest-ip uses all of the next-hops available to BGP that appear in the BGP database.

hash-based — Specifies that the UMH selection is based on hash-based procedures set out in RFC6513, section 5.1.3. The RTM may have just one next-hop to the source, but **hash-based** uses all of the next-hops available to BGP that appear in the BGP database.

tunnel-status — Specifies that UMH selection is based on the state of the tunnel as well as the available unicast routes through the tunnel. Not supported for IPv6.

unicast-rt-pref — When selected, best unicast route will decide which UMH is chosen. All PE routers shall prefer the same route to the UMH for the UMH selection criterion (for example BGP path selection criteria must not influence one PE to choose different UMH from another PE).

vrf-export

Syntax	vrf-export { unicast <i>plcy-or-long-expr</i> [<i>plcy-or-expr</i> [<i>plcy-or-expr</i>]]} no vrf-export
Context	config>service>vprn>mvpn
Description	This command specifies the export policy to control MVPN routes exported from the local VRF to other VRFs on the same or remote PE routers.
Default	vrf-export unicast
Parameters	unicast — Specifies to use unicast VRF export policy for the MVPN. <i>plcy-or-long-expr</i> — The route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long). Allowed values are any string up to 255 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. <i>plcy-or-expr</i> — The route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long). Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. Up to 14 policies can be specified in a single statement.

vrf-import

Syntax	vrf-import { unicast <i>plcy-or-long-expr</i> [<i>plcy-or-expr</i> [<i>plcy-or-expr</i>]]} no vrf-import
Context	config>service>vprn>mvpn
Description	This command specifies the import policy to control MVPN routes imported to the local VRF from other VRFs on the same or remote PE routers.

Default	vrf-import unicast
Parameters	unicast — Specifies to use a unicast VRF import policy for the MVPN. <i>plcy-or-long-expr</i> — The route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long). Allowed values are any string up to 255 characters in length composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. <i>plcy-or-expr</i> — The route policy statement name or a policy logical expression. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. Up to 14 policies can be specified in a single statement.

vrf-target

Syntax	vrf-target {unicast ext-community export unicast ext-community import unicast ext-community} no vrf-target								
Context	config>service>vprn>mvpn								
Description	This command specifies the route target to be added to the advertised routes or compared against the received routes from other VRFs on the same or remote PE routers. vrf-import or vrf-export policies override the vrf-target policy. The no form of the command removes the vrf-target.								
Default	no vrf-target								
Parameters	unicast — Specifies to use unicast vrf-target ext-community for the multicast VPN. <i>ext-comm</i> — 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:{<i>ip-address:comm-val</i> <i>2byte-asnumber:ext-comm-val</i> <i>4byte-asnumber:comm-val</i>} <table> <tr> <td><i>ip-address:</i></td><td>a.b.c.d</td></tr> <tr> <td><i>comm-val:</i></td><td>0 to 65535</td></tr> <tr> <td><i>2byte-asnumber:</i></td><td>1 to 65535</td></tr> <tr> <td><i>4byte-asnumber</i></td><td>0 to 4294967295</td></tr> </table> 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.	<i>ip-address:</i>	a.b.c.d	<i>comm-val:</i>	0 to 65535	<i>2byte-asnumber:</i>	1 to 65535	<i>4byte-asnumber</i>	0 to 4294967295
<i>ip-address:</i>	a.b.c.d								
<i>comm-val:</i>	0 to 65535								
<i>2byte-asnumber:</i>	1 to 65535								
<i>4byte-asnumber</i>	0 to 4294967295								

export

Syntax	export { unicast <i>ext-community</i> }								
Context	config>service>vprn>mvpn>vrf-target								
Description	This command specifies communities to be sent to peers.								
Parameters	unicast — Specifies to use unicast vrf-target ext-community for the multicast VPN. <i>ext-comm</i> — 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:{ <i>ip-address:comm-val</i> <i>2byte-asnumber:ext-comm-val</i> <i>4byte-asnumber:comm-val</i> } <table><tr><td><i>ip-address:</i></td><td>a.b.c.d</td></tr><tr><td><i>comm-val:</i></td><td>0 to 65535</td></tr><tr><td><i>2byte-asnumber:</i></td><td>1 to 65535</td></tr><tr><td><i>4byte-asnumber</i></td><td>0 to 4294967295</td></tr></table>	<i>ip-address:</i>	a.b.c.d	<i>comm-val:</i>	0 to 65535	<i>2byte-asnumber:</i>	1 to 65535	<i>4byte-asnumber</i>	0 to 4294967295
<i>ip-address:</i>	a.b.c.d								
<i>comm-val:</i>	0 to 65535								
<i>2byte-asnumber:</i>	1 to 65535								
<i>4byte-asnumber</i>	0 to 4294967295								

import

Syntax	import { unicast <i>ext-community</i> }								
Context	config>service>vprn>mvpn>vrf-target								
Description	This command specifies communities to be accepted from peers.								
Parameters	unicast — Specifies to use unicast vrf-target ext-community for the multicast VPN. <i>ext-comm</i> — 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:{ <i>ip-address:comm-val</i> <i>2byte-asnumber:ext-comm-val</i> <i>4byte-asnumber:comm-val</i> } <table><tr><td><i>ip-address:</i></td><td>a.b.c.d</td></tr><tr><td><i>comm-val:</i></td><td>0 to 65535</td></tr><tr><td><i>2byte-asnumber:</i></td><td>1 to 65535</td></tr><tr><td><i>4byte-asnumber</i></td><td>0 to 4294967295</td></tr></table>	<i>ip-address:</i>	a.b.c.d	<i>comm-val:</i>	0 to 65535	<i>2byte-asnumber:</i>	1 to 65535	<i>4byte-asnumber</i>	0 to 4294967295
<i>ip-address:</i>	a.b.c.d								
<i>comm-val:</i>	0 to 65535								
<i>2byte-asnumber:</i>	1 to 65535								
<i>4byte-asnumber</i>	0 to 4294967295								

3.8.2.11 Redundant Interface Commands

redundant-interface

Syntax	[no] redundant-interface <i>ip-int-name</i>
Context	config>service>vprn
Description	This command configures a redundant interface.
Parameters	<i>ip-int-name</i> — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

address

Syntax	address { <i>ip-address/mask</i> <i>ip-address netmask</i> } [remote-ip <i>ip-address</i>] no address
Context	config>service>vprn>redundant-interface
Description	This command assigns an IP address mask or netmask and a remote IP address to the interface.
Parameters	<i>ip-address/mask</i> — Assigns an IP address/IP subnet format to the interface. <i>ip-address netmask</i> — 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. Assigns an IP address netmask to the interface. remote-ip <i>ip-address</i> — Assigns a remote IP to the interface.

3.8.2.12 Router Advertisement Commands

router-advertisement

Syntax	[no] router-advertisement
Context	config>service>vprn
Description	This command configures router advertisement properties. By default, it is disabled for all IPv6 enabled interfaces.

The **no** form of the command disables all IPv6 interface. However, the **no interface interface-name** command disables a specific interface.

Default disabled

dns-options

Syntax [no] **dns-options**

Context config>service>vprn>router-advertisement
config>service>vprn>router-advert>interface

Description This command enters the context for configuration of DNS information for Stateless Address Auto-Configuration (SLAAC) hosts.

When specified at the router-advertisement level in the routing context, this command allows configuration of service-wide parameters. These can then be inherited at the interface level by specifying the **config>service>vprn>router-advert>interface>dns-options>include-dns** command.

The **no** form of the command disables configuration of DNS information for Stateless Address Auto-Configuration (SLAAC) hosts.

Default disabled

server

Syntax **server ipv6-address**
no server

Context config>service>vprn>router-advert>dns-options
config>service>vprn>router-advert>interface>dns-options

Description This command specifies the IPv6 DNS servers to include in the RDNSS option in Router Advertisements. When specified at the router advertisement level this applies to all interfaces that have **include-dns** enabled, unless the interfaces have more specific **dns-options** configured.

Parameters *ipv6-address* — Specifies the IPv6 address of the DNS server(s), up to 4 max. Specified as eight 16-bit hexadecimal pieces.

include-dns

Syntax [no] **include-dns**

Context config>service>vprn>router-advert>interface>dns-options

Description	This command enables the Recursive DNS Server (RDNSS) Option in router advertisements. This must be enabled for each interface on which the RDNSS option is required in router advertisement messages. The no form of the command disables the RDNSS option in router advertisements.
Default	disabled

rdnss-lifetime

Syntax	rdnss-lifetime { <i>seconds</i> infinite } no rdnss-lifetime
Context	config>service>vprn>router-advert>dns-options config>service>vprn>router-advert>interface>dns-options
Description	This command specifies the maximum time that the RDNSS address may be used for name resolution by the client. The RDNSS Lifetime must be no more than twice MaxRtrAdvLifetime with a maximum of 3600 seconds.
Default	rdnss-lifetime infinite
Parameters	<i>infinite</i> — Specifies an infinite RDNSS lifetime. <i>seconds</i> — Specifies the time in seconds.
Values	4to 3600

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>service>vprn>router-advertisement
Description	This command configures router advertisement properties on a specific interface. The interface must already exist in the config>router>interface context.
Default	No interfaces are configured by default.
Parameters	<i>ip-int-name</i> — Specifies the interface name. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

current-hop-limit

Syntax	current-hop-limit <i>number</i> no current-hop-limit
Context	config>service>vprn>router-advert>if

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	current-hop-limit 64
Parameters	<i>number</i> — Specifies the hop limit.
Values	0 to 255. A value of zero means there is an unspecified number of hops.

managed-configuration

Syntax	[no] managed-configuration
Context	config>service>vprn>router-advert>if
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. See RFC 3315, <i>Dynamic Host Configuration Protocol (DHCP) for IPv6</i> .
Default	no managed-configuration

max-advertisement-interval

Syntax	[no] max-advertisement-interval <i>seconds</i>
Context	config>service>vprn>router-advert>if
Description	This command configures the maximum interval between sending router advertisement messages.
Default	max-advertisement-interval 600
Parameters	<i>seconds</i> — Specifies the maximum interval in seconds between sending router advertisement messages.
Values	4 to 1800

min-advertisement-interval

Syntax	[no] min-advertisement-interval <i>seconds</i>
Context	config>service>vprn>router-advert>if
Description	This command configures the minimum interval between sending ICMPv6 neighbor discovery router advertisement messages.

Default	min-advertisement-interval 200
Parameters	<i>seconds</i> — Specifies the minimum interval in seconds between sending ICMPv6 neighbor discovery router advertisement messages.
Values	3 to 1350

mtu

Syntax	[no] mtu <i>mtu-bytes</i>
Context	config>service>vprn>router-advert>if
Description	This command configures the MTU for the nodes to use to send packets on the link.
Default	no mtu — The MTU option is not sent in the router advertisement messages.
Parameters	<i>mtu-bytes</i> — Specifies the MTU for the nodes to use to send packets on the link.
Values	1280 to 9212

other-stateful-configuration

Syntax	[no] other-stateful-configuration
Context	config>service>vprn>router-advert>if
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. See RFC 3736, <i>Stateless Dynamic Host Configuration Protocol (DHCP) for IPv6</i> .
Default	no other-stateful-configuration

prefix

Syntax	[no] prefix [<i>ipv6-prefix/prefix-length</i>]
Context	config>service>vprn>router-advert>if
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	<i>ip-prefix</i> — The IP prefix for prefix list entry in dotted decimal notation.
Values	ipv4-prefix a.b.c.d (host bits must be 0)

ipv4-prefix-length	0 to 32
ipv6-prefix	x::x::x::x::x (eight 16-bit pieces) x::x::x::x::d.d.d.d x: [0 to FFFF]H d: [0 to 255]D
ipv6-prefix-length	0 to 128

prefix-length — Specifies a route must match the most significant bits and have a prefix length.

Values 1 to 128

autonomous

Syntax	[no] autonomous
Context	config>service>vprn>router-advert>if>prefix
Description	This command specifies whether the prefix can be used for stateless address autoconfiguration.
Default	autonomous

on-link

Syntax	[no] on-link
Context	config>service>vprn>router-advert>if>prefix
Description	This command specifies whether the prefix can be used for onlink determination.
Default	on-link

preferred-lifetime

Syntax	[no] preferred-lifetime {seconds infinite}
Context	config>service>vprn>router-advert>if
Description	This command configures the remaining length of time in seconds that this prefix will continue to be preferred, such as, time until deprecation. The address generated from a deprecated prefix should not be used as a source address in new communications, but packets received on such an interface are processed as expected.
Default	preferred-lifetime 604800

Parameters	<i>seconds</i> — Specifies the remaining length of time in seconds that this prefix will continue to be preferred. <i>infinite</i> — Specifies that the prefix will always be preferred. A value of 4,294,967,295 represents infinity.

valid-lifetime

Syntax	valid-lifetime { <i>seconds</i> <i>infinite</i> }
Context	config>service>vprn>router-advert>if
Description	This command specifies the length of time in seconds that the prefix is valid for the purpose of on-link determination. A value of all one bits (0xffffffff) represents infinity. The address generated from an invalidated prefix should not appear as the destination or source address of a packet.
Default	valid-lifetime 2592000
Parameters	<i>seconds</i> — Specifies the remaining length of time in seconds that this prefix will continue to be valid. <i>infinite</i> — Specifies that the prefix will always be valid. A value of 4,294,967,295 represents infinity.

reachable-time

Syntax	reachable-time <i>milli-seconds</i> no reachable-time
Context	config>service>vprn>router-advert>if
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	<i>milli-seconds</i> — Specifies the length of time the router should be considered reachable. Values 0 to 3600000

retransmit-time

Syntax	retransmit-timer <i>milli-seconds</i> no retransmit-timer
Context	config>service>vprn>router-advert>if

Description	This command configures the retransmission frequency of neighbor solicitation messages.
Default	no retransmit-time
Parameters	<i>milli-seconds</i> — Specifies how often the retransmission should occur.
Values	0 to 1800000

router-lifetime

Syntax	router-lifetime <i>seconds</i> no router-lifetime
Context	config>service>vprn>router-advert>if
Description	This command sets the router lifetime.
Default	router-lifetime 1800
Parameters	<i>seconds</i> — The length of time, in seconds, (relative to the time the packet is sent) that the prefix is valid for route determination.
Values	0, 4 to 9000 seconds. 0 means that the router is not a default router on this link.

use-virtual-mac

Syntax	[no] use-virtual-mac
Context	config>service>vprn>router-advert>if
Description	This command enables sending router advertisement messages using the VRRP virtual MAC address, provided that the virtual router is currently the master. If the virtual router is not the master, no router advertisement messages are sent. The no form of the command disables sending router advertisement messages.
Default	no use-virtual-mac

3.8.2.13 Network Time Protocol Commands

ntp

Syntax	[no] ntp

Context config>service>vprn

Description This command enters the context to configure Network Time Protocol (NTP) and its operation. It also enables NTP server mode within the VPRN routing instance so that the router will respond to NTP requests from external clients received inside the VPRN.

The **no** form of the command stops the execution of NTP and removes its configuration.

authenticate

Syntax [no] **authenticate**

Context config>service>vprn>ntp

Description This command enables authentication for the NTP server.

authentication-check

Syntax [no] **authentication-check**

Context config>service>vprn>ntp

Description This command provides the option to skip the rejection of NTP PDUs that do not match the authentication key-id, type or key requirements. The default behavior when authentication is configured is to reject all NTP protocol PDUs that have a mismatch in either the authentication key-id, type or key.

When **authentication-check** is enabled, NTP PDUs are authenticated on receipt. However, mismatches cause a counter to be increased, one counter for type and one for key-id, one for type, value mismatches. These counters are visible in a show command.

The **no** form of this command allows authentication mismatches to be accepted; the counters however are maintained.

Default authentication-check — Rejects authentication mismatches.

authentication-key

Syntax **authentication-key** *key-id* {**key** *key*} [**hash** | **hash2**] **type** {**des** | **message-digest**}
no authentication-key *key-id*

Context config>service>vprn>ntp

Description This command sets the authentication key-id, type and key used to authenticate NTP PDUs sent by the broadcast server function toward external clients or to authenticate NTP PDUs received from external unicast clients within the VPRN routing instance. For authentication to work, the authentication key-id, type, and key value must match.

The **no** form of the command removes the authentication key.

- Parameters** *key-id* — Configure the authentication key-id that will be used by the node when transmitting or receiving Network Time Protocol packets.
- Entering the authentication-key command with a key-id value that matches an existing configuration key will result in overriding the existing entry.
- Recipients of the NTP packets must have the same authentication key-id, type, and key value in order to use the data transmitted by this node. This is an optional parameter.
- Values** 1 to 255
- key** — The authentication key associated with the configured key-id, the value configured in this parameter is the actual value used by other network elements to authenticate the NTP packet.
- The key can be any combination of ASCII characters up to 8 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks ("").
- hash** — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.
- hash2** — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.
- type** — This parameter determines if DES or message-digest authentication is used. This is a required parameter; either DES or message-digest must be configured.
- Values** *des* — Specifies that DES authentication is used for this key. The *des* value is not supported in FIPS-140-2 mode.
- message-digest* — Specifies that MD5 authentication in accordance with RFC 2104 is used for this key.

broadcast

- Syntax** **broadcast** {**interface** *ip-int-name*} [**key-id** *key-id*] [**version** *version*] [**ttl** *ttl*]
no broadcast {**interface** *ip-int-name*}
- Context** config>service>vprn>ntp
- Description** This command configures the node to transmit NTP packets on a given interface. Broadcast and multicast messages can easily be spoofed, thus, authentication is strongly recommended.

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

Parameters	<i>ip-int-name</i> — Specifies the local interface on which to transmit NTP broadcast packets. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. Values 32 character maximum key-id <i>key-id</i> — Identifies the configured authentication key and authentication type used by this node to receive and transmit NTP packets to and from an NTP server and peers. If an NTP packet is received by this node both authentication key and authentication type must be valid otherwise the packet will be rejected and an event/trap generated. Values 1 to 255 version <i>version</i> — Specifies the NTP version number that is generated by this node. This parameter does not need to be configured when in client mode in which case all versions will be accepted. Values 2 to 4 Default 4 tll <i>tll</i> — Specifies the IP Time To Live (TTL) value. Values 1 to 255

3.8.2.14 NAT Commands

nat

Syntax	[no] nat
Context	config>service>vprn config>router
Description	This command configures, creates or deletes a NAT instance.

inside

Syntax	inside
Context	config>service>vprn>nat config>router>nat
Description	This command enters the “inside” context to configure the inside NAT instance.

destination-prefix

Syntax	[no] destination-prefix <i>ip-prefix/length</i>
Context	config>service>vprn>nat>inside config>router>nat>inside
Description	This command configures a destination prefix. An (internal) static route will be created for this prefix. All traffic that hits this route will be subject to NAT. The system will not allow a destination-prefix to be configured if the configured nat-policy refers to an IP pool that resides in the same service (as this would result in a routing loop).
Parameters	<i>ip-prefix</i> — Specifies the IP prefix; host bits must be zero (0). Values a.b.c.d <i>length</i> — Specifies the prefix length. Values 0 to 32

dual-stack-lite

Syntax	dual-stack-lite
Context	config>service>vprn>nat config>router>nat>inside
Description	This command enters the context to configure Dual-Stack-Lite NAT parameters.

address

Syntax	[no] address <i>ipv6-address</i>
Context	config>service>vprn>nat>inside>dslite
Description	This command configures a dual-stack-lite IPv6 address The no form of the command removes the value from the configuration.
Parameters	<i>ipv6-address</i> — Specifies the IPv6 address on the interface. Values ipv6-prefix x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x.d.d.d x: [0 to FFFF]H d: [0 to 255]D

tunnel-mtu

Syntax	tunnel-mtu <i>mtu-bytes</i> no tunnel-mtu
Context	config>service>vprn>nat>inside>dslite>address
Description	This command configures the DSLite tunnel MTU for this Dual Stack Lite address. The no form of the command reverts the default.
Default	tunnel-mtu 1500
Parameters	<i>mtu-bytes</i> — Specifies the DSLite tunnel MTU. Values 512 to 9212

subscriber-prefix-length

Syntax	subscriber-prefix-length <i>prefix-length</i> no subscriber-prefix-length
Context	config>service>vprn>nat>inside>dslite
Description	This command configures the IPv6 prefix length of the dual-stack-lite subscribers. The no form of the command reverts the default.
Default	subscriber-prefix-length 128
Parameters	prefix-length <i>prefix-length</i> — Specifies the IPv6 prefix length of the dual-stack-lite subscriber. Values 32 to 64, 128 Default 128

l2-aware

Syntax	l2-aware
Context	config>services>vprn>nat>inside
Description	This command enters the context to configure parameters specific to Layer 2-aware NAT.

address

Syntax	[no] address <i>ip-address/mask</i>

Context	config>services>vprn>nat>inside>l2-aware
Description	This command configures a Layer 2-aware NAT address. This address will act as a local address of the system. Hosts connected to the inside service will be able to ARP for this address. To verify connectivity, a host can also ping the address. This address is typically used as next hop of the default route of a Layer 2-aware host. The given mask defines a Layer 2-aware subnet. The (inside) IP address used by a Layer 2-aware host must match one of the subnets defined here or it will be rejected.
Parameters	<i>ip-address</i> — Specifies the IP address in a.b.c.d format. <i>mask</i> — Specifies the mask.
Values	16 to 32

nat-policy

Syntax	nat-policy <i>nat-policy-name</i> no nat-policy
Context	config>services>vprn>nat>inside config>router>nat>inside
Description	This command configures the NAT policy that will be used for large-scale NAT in this service.
Parameters	<i>nat-policy-name</i> — Specifies the NAT policy name.
Values	32 chars max

redundancy

Syntax	redundancy
Context	config>service>vprn>nat>inside config>service>vprn>nat>outside>pool
Description	This command enters the context to configure redundancy parameters.

peer

Syntax	peer <i>ip-address</i> no peer
Context	config>service>vprn>nat>inside>redundancy
Description	This command configures the IP address of the NAT redundancy peer in the realm of this virtual router instance.

steering-route

Syntax	steering-route <i>ip-prefix/length</i> no steering-route
Context	config>service>vprn>nat>inside>redundancy
Description	This command configures specifies the IP address and prefix length of the steering route. The steering route is used in the realm of this virtual router instance as an indirect next-hop for all the traffic that must be routed to the large scale NAT function.

outside

Syntax	outside
Context	config>service>vprn>nat config>router>nat
Description	This command enters the “outside” context to configure the outside NAT instance.

pool

Syntax	pool <i>nat-pool-name</i> [nat-group <i>nat-group-id</i> type <i>pool-type</i> [no-allocate] [create] no pool <i>nat-pool-name</i>
Context	config>service>vprn>nat>outside config>router>nat>outside
Description	This command configures a NAT pool.
Parameters	<i>nat-pool-name</i> — Specifies the NAT pool name. Values 32 chars max <i>nat-group-id</i> — Specifies the NAT group ID. Values 1 to 4 create — This parameter must be specified to create the instance. <i>pool-type</i> — Specifies the pool type, either large-scale or L2-aware.

address-range

Syntax	address-range <i>start-ip-address end-ip-address</i> [create] no address-range <i>start-ip-address end-ip-address</i>
Context	config>service>vprn>nat>outside>pool

config>router>nat>outside>pool

- Description** This command configures a NAT address range.
- Parameters** *start-ip-address* — Specifies the beginning IP address in a.b.c.d form.
end-ip-address — Specifies the ending IP address in a.b.c.d. form.
create — This parameter must be specified to create the instance.

description

- Syntax** **description** *description-string*
no description
- Context** config>service>vprn>nat>outside>pool>address-range
config>service>vprn>nat>outside>pool
config>router>nat>outside>pool>address-range
config>router>nat>outside>pool
- Description** This command configures the description for the NAT address range.
- Parameters** *description-string* — Specifies the NAT address range description.
Values 80 chars max

drain

- Syntax** [no] **drain**
- Context** config>service>vprn>nat>outside>pool>address-range
config>router>nat>outside>pool>address-range
- Description** This command starts or stops draining this NAT address range. When an address-range is being drained, it will not be used to serve new hosts. Existing hosts, however, will still be able to use the address that was assigned to them even if it is being drained. An address-range can only be deleted if the parent pool is shut down or if the range itself is effectively drained (no hosts are using the addresses anymore).

mode

- Syntax** **mode** {**auto** | **napt**}
no mode
- Context** config>service>vprn>nat>outside>pool
- Description** This command configures the mode of operation of this NAT address pool.

The mode value is only relevant while the value of pool type is equal to largeScale; while the value of pool type is equal to l2Aware, the mode of operation is always NATP.

port-forwarding-range

Syntax	port-forwarding-range <i>range-end</i> no port-forwarding-range
Context	config>service>vprn>nat>outside>pool
Description	This command configures the end of the port range available for port forwarding. The start of the range is always equal to one. The actual maximum value of the range end may be restricted to less than 65535 depending on the value of the objects port reservation type and port reservation value and on system specifications.
Default	port-forwarding-range 1023
Parameters	<i>range-end</i> — Specifies the mode of operation of this NAT pool. Values 1023 to 65535

port-reservation

Syntax	port-reservation blocks <i>num-blocks</i> port-reservation ports <i>num-ports</i> no port-reservation
Context	config>service>vprn>nat>outside>pool config>router>nat>outside>pool
Description	This command configures the size of the port-block that will be assigned to a host that is served by this pool. The number of ports configured here will be available to UDP, TCP and ICMP (as identifiers).
Parameters	<i>num-blocks</i> — Specifies the number of port-blocks per IP address. Setting this parameter to one (1) for large scale NAT will enable 1:1 NAT for IP addresses in this pool. Values 1 to 64512 <i>num-ports</i> — Specifies the number of ports per block. Values 1 to 32256

export

Syntax	export ip-prefix/length no export						
Context	config>service>vprn>nat>outside>pool>redundancy						
Description	This command installs the export route in the routing table for active NAT pools. Once the export route is in the routing table, it can be advertised in the network via a routing protocol. NAT pools in the standby or disabled state will not advertise the export route. A NAT pool becomes active when it becomes operationally UP, and there is no monitoring route (which is also the export route from the peer) present in the routing node (as received from the network). The pool will transition into standby state in case that the monitoring route (or export route from the peer) is already present in the routing table. In other words, the monitoring route is already advertised as an export route from the peering node with active NAT pool. The export route can be advertised only from: • The active lead pool. • Active pool for which fate-sharing is disabled.						
Default	no export						
Parameters	ip-prefix/length — Specifies the IP prefix and length.						
	Syntax:						
	<table><tr><td>ip-prefix/length:</td><td>ip-prefix</td><td>a.b.c.d</td></tr><tr><td></td><td>ip-prefix-length</td><td>0 to 32</td></tr></table>	ip-prefix/length:	ip-prefix	a.b.c.d		ip-prefix-length	0 to 32
ip-prefix/length:	ip-prefix	a.b.c.d					
	ip-prefix-length	0 to 32					
	Values	0, 4, 16					

follow

Syntax	follow router <i>router-instance</i> pool name no follow
Context	config>service>vprn>nat>outside>pool>redundancy config>router> nat>outside>pool>redundancy
Description	This command implicitly enables Pool Fate-Sharing Group (PFSG) which is required in case of multiple NAT policies per inside routing context. A NAT pool configured with this command will not advertise or monitor any route in order to change its (activity) state but instead it will directly follow the state of the lead pool in the PFSG. Once the lead pool changes its (activity) state, all the remaining pools following the lead pool will change their state accordingly.

Default	no follow
Parameters	router <i>router-instance</i> — Specifies the routing instance where the lead pool resides.
	Values <router-name> <service-id>
	router-name - Base
	service-id - 1 to 2147483647
	pool <i>name</i> — Specifies the pool whose activity state is being shared up to 32 characters in length.

monitor

Syntax	monitor <i>ip-prefix/length</i> no monitor						
Context	config>service>vprn>nat>outside>pool>redundancy config>router> nat>outside>pool>redundancy						
Description	This command configures the monitoring route based on which the NAT multi-chassis switchover is triggered. Monitoring route of a NAT pool on the local node must match the export route of a corresponding NAT pool on the peering node. Presence of the monitoring route in the routing table is an indication that the peering NAT pool is active (since it is advertising its export route). The disappearance of the monitoring route from the routing table is an indication that the peering pool has failed and consequently the nodal switchover is triggered, the local pool becomes active and its export route is consequently advertised. The export route can be advertised only from: • The active lead pool. • Active pool for which fate-sharing is disabled.						
Parameters	<i>ip-prefix/length</i> — Specifies the IP prefix and length.						
	Syntax:						
	<table><tr><td>ip-prefix/length:</td><td>ip-prefix</td><td>a.b.c.d</td></tr><tr><td></td><td>ip-prefix-length</td><td>0 to 32</td></tr></table>	ip-prefix/length:	ip-prefix	a.b.c.d		ip-prefix-length	0 to 32
ip-prefix/length:	ip-prefix	a.b.c.d					
	ip-prefix-length	0 to 32					

subscriber-limit

Syntax	subscriber-limit <i>limit</i> no subscriber-limit
Context	config>service>vprn>nat>outside>pool
Description	This command configures the maximum number of subscribers per outside IP address.

If multiple port blocks per subscriber are used, the block size is typically small; all blocks assigned to a given subscriber belong to the same IP address; the subscriber limit guarantees that any subscriber can get a minimum number of ports.

Parameters *limit* — Specifies the maximum number of subscribers per outside IP address.
Values 1 to 65535

watermarks

Syntax **watermarks high** *percentage-high* **low** *percentage-low*
no watermarks

Context config>service>vprn>nat>outside>pool
config>router>nat>outside>pool

Description This command configures the watermarks for this NAT pool.

Parameters *percentage-high* — Specifies the high percentage.
Values 2 to 100
percentage-low — Specifies the low percentage.
Values 1 to 99

3.8.2.15 Subscriber Interface Commands

subscriber-interface

Syntax **subscriber-interface** *ip-int-name* [**fwd-service** *service-id* **fwd-subscriber-interface** *ip-int-name*]
no subscriber-interface *ip-int-name*

Context config>service>vprn

Description This command allows the operator to create special subscriber-based interfaces. It is used to contain multiple group interfaces. Multiple subnets associated with the subscriber interface can be applied to any of the contained group interfaces in any combination. The subscriber interface allows subnet sharing between group interfaces.

Use the **no** form of the command to remove the subscriber interface.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

fwd-service *service-id* — Specifies the forwarding service ID for a subscriber interface in a retailer context.

fwd-subscriber-interface *ip-int-name* — Specifies the forwarding subscriber interface for a subscriber interface in a retailer context.

address

Syntax	[no] address { <i>ip-address/mask</i> <i>ip-address netmask</i> } [gw-ip-address <i>ip-address</i>] [populate-host-routes] [track-srrp <i>srrp-instance</i>] [holdup-time <i>msecs</i>]
Context	config>service>vprn>subscriber-interface
Description	This command configures the local subscriber subnets available on a subscriber IP interface. The configured <i>ip-address</i> and mask define the address space associated with the subscriber subnet. Up to 16 IP subnets can be created on a single subscriber IP interface. Each subnet supports a locally owned IP host address within the subnet that is not expected to appear on other routers that may be servicing the same subscriber subnet. For redundancy purposes, the keyword gw-address defines a separate IP address within the subnet for Subscriber Routed Redundancy Protocol (SRRP) routing. This IP address must be the same on the local and remote routers participating in a common SRRP instance. In SRRP, a single SRRP instance is tied to a group IP interface. The group IP interface is contained directly within a subscriber IP interface context and thus directly associated with the subscriber subnets on the subscriber IP interface. The SRRP instance is also indirectly associated with any subscriber subnets tied to the subscriber interface through wholesale/retail VPRN configurations. With the directly-associated and the indirectly-associated subscriber interface subnets, a single SRRP instance can manage hundreds of SRRP gateway IP addresses. This automatic subnet association to the SRRP instance is different from VRRP where the redundant IP address is defined within the VRRP context. Defining an SRRP gateway IP address on a subscriber subnet is not optional when the subnet is associated with a group IP interface with SRRP enabled. Enabling SRRP (no shutdown) fails if one or more subscriber subnets do not have an SRRP gateway IP address defined. Creating a new subscriber subnet without an SRRP gateway IP address defined fails when the subscriber subnet is associated with a group IP interface with an active SRRP instance. Once SRRP is enabled on a group interface, the SRRP instance will manage the ARP response and routing behavior for all subscriber hosts reachable through the group IP interface. The no form of the command removes the address from a subscriber subnet. The address command for the specific subscriber subnet must be executed without the <i>gw-address</i> parameter. To succeed, all SRRP instances associated with the subscriber subnet must be removed or shutdown.
Parameters	<i>ip-address/mask</i> <i>ip-address netmask</i> — Specifies the address space associated with the subscriber subnet.

gw-ip-address *ip-address* — Specifies a separate IP address within the subnet for SRRP routing purposes. This parameter must be followed by a valid IP interface that exists within the subscriber subnet created by the address command. The defined gateway IP address cannot currently exist as a subscriber host (static or dynamic). If the defined ip-address already exists as a subscriber host address, the address command fails. The specified ip-address must be unique within the system.

The gw-address parameter may be specified at anytime. If the subscriber subnet was created previously, executing the address command with a gw-address parameter will simply add the SRRP gateway IP address to the existing subnet.

If the address command is executed without the gw-address parameter when the subscriber subnet is associated with an active SRRP instance, the address fails. If the SRRP instance is inactive or removed, executing the address command without the gw-address parameter will remove the SRRP gateway IP address from the specified subscriber subnet.

If the address command is executed with a new gw-address, all SRRP instances currently associated with the specified subscriber subnet will be updated with the new SRRP gateway IP address.

populate-host-routes — Specifies to populate subscriber-host routes in local FIB. Storing them in FIB benefits topologies only where the external router advertises more specific routes than the one corresponding to locally configured subscriber-interface subnets.

allow-unmatching-subnets

Syntax	[no] allow-unmatching-subnets
Context	config>service>vprn>subscriber-interface
Description	This command specifies whether subscriber hosts with a subnet that does not match any of the subnets configured on this interface, are allowed.

group-interface

Syntax	[no] group-interface <i>ip-int-name</i>
Context	config>service>vprn>subscriber-interface
Description	This command enters the context to configure a group interface. A group interface is an interface that may contain one or more SAPs. This interface is used in triple-play services where multiple SAPs are part of the same subnet.
Parameters	<i>ip-int-name</i> — Configures the interface group name. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

arp-host

Syntax	arp-host
Context	config>service>vprn>sub-if>grp-if
Description	This command enters the context to configure ARP host parameters.

host-limit

Syntax	host-limit <i>max-num-hosts</i> no host-limit
Context	config>service>vprn>sub-if>grp-if>arp-host
Description	This command configures the maximum number of ARP hosts.
Parameters	<i>max-num-hosts</i> — Specifies the maximum number of ARP hosts. Values 1 to 32767

min-auth-interval

Syntax	min-auth-interval <i>min-auth-interval</i> no min-auth-interval
Context	config>service>vprn>sub-if>grp-if>arp-host
Description	This command configures the minimum authentication interval.
Parameters	<i>min-auth-interval</i> — Specifies the minimum authentication interval. Values 1 to 6000

sap-host-limit

Syntax	sap-host-limit <i>max-num-hosts-sap</i> no sap-host-limit
Context	config>service>vprn>sub-if>grp-if>arp-host
Description	This command configures the maximum number of ARP hosts per SAP.
Parameters	<i>max-num-hosts-sap</i> — Specifies the maximum number of ARP hosts per SAP allowed on this IES interface. Values 1 to 32767

3.8.2.16 Interface Commands

interface

Syntax	interface <i>ip-int-name</i> no interface <i>ip-int-name</i>
Context	config>service>vprn
Description	This command creates a logical IP routing interface for a Virtual Private Routed Network (VPRN). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.

The **interface** command, under the context of services, is used to create and maintain IP routing interfaces within VPRN service IDs. The **interface** command can be executed in the context of an VPRN service ID. The IP interface created is associated with the service core network routing instance and default routing table. The typical use for IP interfaces created in this manner is for subscriber Internet access.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config router interface** and **config service vprn interface**. Interface names must not be in the dotted decimal notation of an IP address. For example, the name "1.1.1.1" is not allowed, but "int-1.1.1.1" is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

The available IP address space for local subnets and routes is controlled with the **config router service-prefix** command. The **service-prefix** command administers the allowed subnets that can be defined on service IP interfaces. It also controls the prefixes that may be learned or statically defined with the service IP interface as the egress interface. This allows segmenting the IP address space into **config router** and **config service** domains.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All VPRN IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

The **no** form of this command removes the interface and all the associated configuration. The interface must be administratively shutdown before issuing the **no interface** command.

For VPRN services, the IP interface must be shutdown before the SAP on that interface may be removed. VPRN services do not have the **shutdown** command in the SAP CLI context. VPRN service SAPs rely on the interface status to enable and disable them.

Parameters	<i>ip-int-name</i> — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for config router interface and config service vprn interface commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. If <i>ip-int-name</i> already exists within the service ID, the context will be changed to maintain that IP interface. If <i>ip-int-name</i> 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 <i>ip-int-name</i> does not exist, the interface is created and context is changed to that interface for further command processing.

active-cpm-protocols

Syntax	[no] active-cpm-protocols
Context	config>service>vprn>if
Description	This command enables CPM protocols on this interface.

address

Syntax	address {ip-address/mask ip-address netmask} [broadcast all-ones host-ones] [track-srrp srrp-instance] no address
Context	config>service>vprn>if config>service>vprn>nw-if
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. Use the secondary command to assign multiple addresses. 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 7750 SR. The local subnet that the address command defines must be part of the services address space within the routing context using the config router service-prefix command. The default is to disallow the complete address space to services. Once a portion of the address space is allocated as a service prefix, that portion can be made unavailable for IP interfaces defined within the config router interface CLI context for network core connectivity with the exclude option in the config router service-prefix command.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. The show commands display CIDR notation and is stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

Use the **no** form of this command to remove the IP address assignment from the IP interface. When the **no address** command is entered, the interface becomes operationally down.

Address	Admin state	Oper state
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an administratively up state and an address is assigned, it becomes operationally up and the protocol interfaces and the MPLS LSPs associated with that IP interface will be reinitialized.

- Parameters**
- ip-address* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).
 - /* — The forward slash is a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the “/” and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted decimal mask must follow the prefix.
 - mask-length* — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 0 – 30. A mask length of 32 is reserved for system IP addresses.
 - mask* — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical ‘AND’ function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. A mask of 255.255.255.255 is reserved for system IP addresses.

broadcast — The optional **broadcast** parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which 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 — The **all-ones** keyword following the **broadcast** parameter specifies the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.

host-ones — The **host-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default broadcast address used by an IP interface.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

track-srrp — Specifies the SRRP instance ID that this interface route needs to track.

allow-directed-broadcasts

Syntax	[no] allow-directed-broadcasts
Context	config>service>vprn>if config>service>vprn>nw-if
Description	This command controls the forwarding of directed broadcasts out of the IP interface. A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address on another IP interface. The allow-directed-broadcasts command on an IP interface enables or disables the transmission of packets destined to the subnet broadcast address of the egress IP interface. When enabled, a frame destined to the local subnet on this IP interface will be sent as a subnet broadcast out this interface. Care should be exercised when allowing directed broadcasts as it is a well-known mechanism used for denial-of-service attacks. When disabled, directed broadcast packets discarded at this egress IP interface will be counted in the normal discard counters for the egress SAP.

By default, directed broadcasts are not allowed and will be discarded at this egress IP interface.

The **no** form of this command disables the forwarding of directed broadcasts out of the IP interface.

Default no allow-directed-broadcasts — Directed broadcasts are dropped.

arp-learn-unsolicited

Syntax [no] **arp-learn-unsolicited**

Context config>service>vprn>if

Description This command allows the ARP application to learn new entries based on any received ARP message (GARP/ARP-Request/ARP-Reply, such as any frame with ethertype 0x0806).

The **no** form of this command disables the above behavior and ARP entries are only learned when needed, that is, when the router receives an ARP-reply after an ARP-request triggered by some received traffic.

arp-populate-host-route

Syntax [no] **arp-populate-host-route**

Context config>service>vprn>if

Description This command enables the addition or deletion of host routes in the route table derived from ARP entries in the ARP cache. In order to enable this command, the interface must be shutdown. The command triggers the population of host routes in the route table out of their corresponding static, dynamic, or EVPN types in the ARP table. ARP entries installed by subscriber management, local interfaces, and others, do not create host routes.

The **no** form of this command disables the creation of host routes from the ARP cache.

arp-proactive-refresh

Syntax [no] **arp-proactive-refresh**

Context config>service>vprn>if

Description This command enables the router to always send out a refresh message 30 seconds prior to the timeout of the entry (a single refresh message with no retries).

The **no** form of this command sets the default behavior, in which an entry is marked as stale 30 seconds prior to age-out, and the router only sends an ARP request to refresh the entry if the IOM receives traffic that uses it. If so, the IOM asks the ARP application to send a refresh message. With **arp-proactive-refresh** enabled, the ARP module sends a refresh message regardless of the IOM receiving traffic.

arp-route-tag

Syntax	arp-route-tag <i>tag</i> arp-route-tag
Context	config>service>vprn>if
Description	This command adds a route tag to the ARP-ND host routes generated from the ARP entries in the interface. As any other route tag, it can be used to match ARP-ND routes in BGP export policies. The no form of this command removes the route tag for the ARP-ND host routes.
Parameters	<i>tag</i> — Specifies the route tag value. Values 1 to 255

bfd

Syntax	bfd <i>transmit-interval</i> [receive <i>receive-interval</i>] [multiplier <i>multiplier</i>] [echo-receive <i>echo-interval</i>] [type <i>cpm-np</i>] no bfd
Context	config>service>vprn>if config>service>vprn>if>ipv6 config>service>vprn>nw-if
Description	This command specifies the BFD parameters for the associated IP interface. If no parameters are defined the default value are used. The multiplier specifies the number of consecutive BFD messages that must be missed from the peer before the BFD session state is changed to down and the upper level protocols (OSPF, IS-IS, BGP or PIM) is notified of the fault. The no form of the command removes BFD from the associated IGP protocol adjacency.



Note: On the 7750 SR and 7950 XRS, the *transmit-interval*, **receive** *receive-interval*, and **echo-receive** *echo-interval* values can only be modified to a value less than 100 when:

1. The **type cpm-np** option is explicitly configured.
2. The interval is specified 10 to 100000.
3. The service is shut down (**shutdown**)
4. The service is re-enabled (**no shutdown**)

To remove the **type cpm-np** option, re-issue the **bfd** command without specifying the **type** parameter.

Default	no bfd
Parameters	<i>transmit-interval</i> — Sets the transmit interval for the BFD session. Values 10 to 100000 (see the Note above) Default 100 <i>receive receive-interval</i> — Sets the receive interval for the BFD session. Values 10 to 100000 (see the Note above) Default 100 <i>multiplier multiplier</i> — Sets the multiplier for the BFD session. Values 3 to 20 Default 3 echo-receive echo-interval — Sets the minimum echo receive interval, in milliseconds, for the BFD session. Values 100 to 100000 10 to 100000 (applies to the 7750 SR; see the Note above) Default 100 type cpm-np — Specifies that BFD sessions associated with this interface will be created on the CPM network processor to allow for fast timers down to 10 ms granularity.

cflowd-parameters

Syntax	cflowd-parameters no cflowd-parameters
Context	config>service>vprn>if
Description	This command creates the configuration context to configure cflowd parameters for the associated IP interfaces.

cflowd is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement.

At a minimum, the **sampling** command must be configured within this context in order to enable cflowd sampling, otherwise traffic sampling will not occur.

Default no cflowd-parameters

sampling

Syntax	sampling {unicast multicast} type {acl interface} [direction {ingress-only egress-only both}] no sampling {unicast multicast}
Context	config>service>vprn>if>cflowd-parameters config>service>vprn>nw-if>cflowd-parameters config>service>vprn>sub-if>grp-if>cflowd-parameters
Description	This command enables and configures the cflowd sampling behavior to collect traffic flow samples through a router for analysis. This command can be used to configure the sampling parameters for unicast and multicast traffic separately. If sampling is not configured for either unicast or multicast traffic, then that type of traffic will not be sampled. If cflowd is enabled without either egress-only or both specified or with the ingress-only keyword specified, then only ingress sampling will be enabled on the associated IP interface. The no form of the command disables the associated type of traffic sampling on the associated interface.
Default	no sampling
Parameters	unicast — Specifies that the sampling command will control the sampling of unicast traffic on the associated interface/SAP. multicast — Specifies that the sampling command will control the sampling of multicast traffic on the associated interface/SAP. type — Specifies the sampling type. Values acl — Specifies that the sampled traffic is controlled via an IP traffic filter entry with the action “filter-sample” configured. interface — Specifies that all traffic entering or exiting the interface is subject to sampling. direction — Specifies the direction in which to collect traffic flow samples. Values ingress-only — Enables ingress sampling only on the associated interface.

egress-only — Enables egress sampling only on the associated interface.

both — Enables both ingress and egress cflowd sampling.

cpu-protection

Syntax	cpu-protection <i>policy-id</i> no cpu-protection
Context	config>service>vprn>if config>service>vprn>nw-if
Description	This command assigns an existing CPU protection policy to the associated service interface. For these interface types, the per-source rate limit is not applicable. The CPU protection policies are configured in the config>sys>security>cpu-protection>policy <i>cpu-protection-policy-id</i> context. If no CPU protection policy is assigned to a service interface, then a the default policy is used to limit the overall-rate. The no form of the command removes CPU protection policy association from the interface, resulting in no default rate limiting of control packets.
Default	cpu-protection 254 (for access interfaces) cpu-protection 255 (for network interfaces) no cpu-protection (for video interfaces)
Parameters	<i>policy-id</i> — Specifies an existing CPU protection policy. Values 1 to 255

cpu-protection

Syntax	cpu-protection <i>policy-id</i> [mac-monitoring] no cpu-protection
Context	config>service>vprn>if config>service>vprn>if>sap
Description	This command assigns an existing CPU protection policy to the associated service group interface SAP, interface or MSAP policy. The CPU protection policies are configured in the config>sys>security>cpu-protection>policy <i>cpu-protection-policy-id</i> context. If no CPU protection policy is assigned to a service group interface SAP, then a the default policy is used to limit the overall-rate.

Default	cpu-protection 254 (for access interfaces) cpu-protection 255 (for network interfaces) no cpu-protection (for video interfaces) The configuration of no cpu-protection returns the interface/SAP to the default policies as shown above.
Parameters	<i>policy-id</i> — Specifies an existing CPU protection policy. Values 1 to 255 mac-monitoring — When specified, the per MAC rate limiting should be performed, using the per-source-rate from the associated cpu-protection policy.

dad-disable

Syntax	[no] dad-disable
Context	config>service>vprn>if>ipv6
Description	This command disables duplicate address detection (DAD) on a per-interface basis. This prevents the router from performing a DAD check on the interface. All IPv6 addresses of an interface with DAD disabled, immediately enter a preferred state, without checking for uniqueness on the interface. This is useful for interfaces which enter a looped state during troubleshooting and operationally disable themselves when the loop is detected, requiring manual intervention to clear the DAD violation. The no form of the command turns off dad-disable on the interface.
Default	no dad-disable

dist-cpu-protection

Syntax	dist-cpu-protection <i>policy-name</i> no dist-cpu-protection
Context	config>service>vprn>nw-if
Description	This command assigns a Distributed CPU Protection (DCP) policy to the network interface. Only a valid created DCP policy can be assigned to a network interface (this rule does not apply to templates such as an msap-policy).
Default	If no dist-cpu-protection policy is assigned to the VPRN network interface, then the default network DCP policy (_default-network-policy) is used. If no DCP functionality is required on the VPRN network interface then an empty DCP policy can be created and explicitly assigned to the VPRN network interface.

Parameters *policy-name* — Specifies the name of the DCP policy up to 32 characters in length

dist-cpu-protection

Syntax **dist-cpu-protection** *policy-name*
no dist-cpu-protection

Context config>service>vprn>if>sap
 config>service>vprn>sub-if>grp-if>sap

Description This command assigns a Distributed CPU Protection (DCP) policy to the network interface. Only a valid created DCP policy can be assigned to a SAP (this rule does not apply to templates such as an msap-policy).

Default If no dist-cpu-protection policy is assigned to a SAP, then the default access DCP policy (_default-access-policy) is used.

 If no DCP functionality is required on the SAP, then an empty DCP policy can be created and explicitly assigned to the SAP.

Parameters *policy-name* — Specifies the name of the DCP policy up to 32 characters in length

if-attribute

Syntax **if-attribute**

Context config>service>vprn>if

Description This command creates the context to configure or apply IP interface attributes such as administrative group (admin-group) or Shared Risk Loss Group (SRLG).

admin-group

Syntax **admin-group** *group-name* [*group-name*]
no admin-group *group-name* [*group-name*]
no admin-group

Context config>service>vprn>if>if-attribute

Description This command configures the admin group membership of an interface. The user can apply admin groups to an IES, VPRN, network IP, or MPLS interface. Once an admin group is bound to one or more interface, its value cannot be changed until all bindings are removed.

 The configured admin-group membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.

Only the admin groups bound to an MPLS interface are advertised area-wide in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

The **no** form of this command deletes one or more of the admin-group memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.

Parameters *group-name* — Specifies the name of the group with up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain. Each single operation of the **admin-group** command allows a maximum of 5 groups to be specified at a time. A maximum of 32 groups can be added to a given interface through multiple operations.

srlg-group

Syntax **srlg-group** *group-name* [*group-name*]
no srlg-group *group-name* [*group-name*]
no srlg-group

Context config>service>vprn>if>if-attribute

Description This command configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface.

Once an SRLG group is bound to one or more interface, its value cannot be changed until all bindings are removed.

The configured SRLG membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.

Only the SRLGs bound to an MPLS interface are advertised area-wide in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

The **no** form of this command deletes one or more of the SRLG memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.

Parameters *group-name* — Specifies the name of the group, up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain. An interface can belong to up to 64 SRLG groups. Each single operation of the **srlg-group** command allows a maximum of 5 groups to be specified at a time.

ingress

Syntax **ingress**

Context	config>service>vprn>if
Description	This command enters context to configure ingress parameters for network interfaces.

policy-accounting

Syntax	policy-accounting <i><template-name></i> no policy-accounting
Context	config>service>vprn>if>ingress
Description	This command configures the service VPRN interface ingress policy accounting
Parameters	<i>template-name</i> — Specifies the template name, up to 32 characters.

ip-helper-address

Syntax	ip-helper-address <i>gateway-address</i> no ip-helper-address
Context	config>service>vprn>if
Description	This command enables broadcast UDP packets received on the associated interface to be redirected to the specified gateway address and then forwarded on to the gateway. The no form of this command removes the gateway address from the interface configuration and stops the UDP broadcast redirect function.
Parameters	<i>gateway-address</i> — Specifies the IPv4 address of the target UDP broadcast gateway.

ip-mtu

Syntax	ip-mtu <i>octets</i> no ip-mtu
Context	config>service>vprn>if
Description	This command configures the IP maximum transmit unit (packet) for this interface. The no form of the command returns the default value.
Default	no ip-mtu

ipcp

Syntax	ipcp
Context	config>service>vprn>if
Description	This command creates allows access to the IPCP context within the interface configuration. Within this context, IPCP extensions can be configured to define such things as the remote IP address and DNS IP address to be signaled via IPCP on the associated PPP interface. This command is only applicable if the associated SAP/port is a PPP/MLPPP interface.

dns

Syntax	dns ip-address [secondary ip-address] dns secondary ip-address no dns [ip-address] [secondary ip-address]
Context	config>service>vprn>if>ipcp
Description	This command defines the dns address(es) to be assigned to the far-end of the associated PPP/MLPPP link via IPCP extensions. This command is only applicable if the associated SAP/port is a PPP/MLPPP interface with an IPCP encapsulation. The no form of the command deletes either the specified primary DNS address, secondary DNS address or both addresses from the IPCP extension peer-ip-address configuration.
Default	no dns
Parameters	<i>ip-address</i> — Specifies a unicast IPv4 address for the primary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions. secondary ip-address — Specifies a unicast IPv4 address for the secondary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions.

peer-ip-address

Syntax	peer-ip-address ip-address no peer-ip-address
Context	config>service>vprn>if>ipcp
Description	This command defines the remote IP address to be assigned to the far-end of the associated PPP/MLPPP link via IPCP extensions.

This command is only applicable if the associated SAP/port is a PPP/MLPPP interface with an IPCP encapsulation.

The interface must be shut down to modify the IPCP configuration.

The **no** form of the command deletes the IPCP extension peer-ip-address configuration.

Default	no peer-ip-address (0.0.0.0)
Parameters	<i>ip-address</i> — Specifies a unicast IPv4 address to be signaled to the far-end of the associated PPP/MLPPP link by IPCP extensions.

ipv6

Syntax	[no] ipv6
Context	config>service>vprn>if
Description	This command configures an IPv6 interface.

address

Syntax	address <i>ipv6-address/mask</i> [eui-64] [track-srrp <i>srrp-instance</i>] [modifier <i>cga-modifier</i>] [dad-disable] [primary-preference <i>primary-preference</i>] no address <i>ipv6-address/prefix-length</i>
Context	config>service>vprn>if>ipv6
Description	This command assigns an IPv6 address to the interface. Up to 16 total primary and secondary IPv4 and IPv6 addresses can be assigned to network interfaces, and up to 256 to access interfaces.



Caution: Configurations must not exceed 16 secondary IP addresses when IPsec, GRE, L2TPv3, or IP in IP protocols are active on an access interface.

Parameters	<i>ipv6-address/prefix-length</i> — Specifies the IPv6 address on the interface.

Values

ipv6-address/prefix:	ipv6-address	x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:d.d.d.d
		x [0 to FFFF]H
		d [0 to 255]D
prefix-length		1 to 128

eui-64 — When the **eui-64** keyword is specified, a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from MAC address on Ethernet interfaces. For interfaces without a MAC address, for example ATM interfaces, the Base MAC address of the chassis is used.

track-srrp *srrp-instance* — Specifies the SRRP instance ID that this interface route needs to track.

Values 1 to 4294967295

modifier *cga-modifier* — Specifies the modifier in 32 hexadecimal nibbles.

Values 0x0–0xFFFFFFFF

dad-disable — Disables Duplicate Address Detection (DAD) and sets the address to preferred, even if there is a duplicated address.

primary-preference *primary-preference* — Specifies a *primary-preference* index to an IPv6 address of the interface to enforce the order in which the address is used by control plane protocols and applications which require a fixed address of the interface. These include LDP and Segment Routing.

When originating packets from this interface, the source IPv6 address follows the selection rules in RFC 6724 except for the specific cases where a fixed address is required. In the latter case, the IPv6 address with the lowest primary-preference index is selected. If the selected address is removed, the system selects the IPv6 address with the next lowest *primary-preference* index.

The system assigns the next available index value to any IPv6 address of the interface when configured without the *primary-preference* index value specified. The address index space is unique across all addresses of a given interface.

Values 1 to 4294967295

dhcp6-relay

Syntax	[no] dhcp6-relay
Context	config>service>vprn>if>ipv6
Description	This command configures DHCPv6 relay parameters for the VPRN interface.

dhcp6-server

Syntax	[no] dhcp6-server
Context	config>service>vprn>if>ipv6
Description	This command configures DHCPv6 server parameters for the VPRN interface.

icmp6

Syntax	icmp6
Context	config>service>vprn>if>ipv6
Description	This command configures ICMPv6 for the interface.

link-local-address

Syntax	link-local-address <i>ipv6-address</i> [dad-disable] no link-local-address
Context	config>router>if>ipv6 config>service>vprn>if>ipv6 config>service>vprn>if>ipv6
Description	This command configures the IPv6 link local address. The no form of the command removes the configured link local address, and the router automatically generates a default link local address.



Caution: Removing a manually configured link local address may impact routing protocols or static routes that have a dependency on that address. It is not recommended to remove a link local address when there are active IPv6 subscriber hosts on an IES or VPRN interface.

Parameters	dad-disable — Disables Duplicate Address Detection (DAD) and sets the address to preferred, even if there is a duplicated address.

local-proxy-nd

Syntax	[no] local-proxy-nd
Context	config>service>vprn>if>ipv6
Description	This command enables or disables neighbor discovery on the interface.

nd-learn-unsolicited

Syntax	nd-learn-unsolicited { global link-local both } no nd-learn-unsolicited
Context	config>service>vprn>if>ipv6

Description	This command enables the ability to learn neighbor entries out of received unsolicited Neighbor Advertisement messages, with or without the solicited flag set. The command can be enabled for global addresses, link-local addresses, or for both. The no form of the command makes the router follow standard RFC 4861 behavior for learning of neighbor entries. • If an unsolicited NA (regardless of the S flag) is received from a neighbor that is not yet in the ND cache, the NA is ignored in line with RFC 4861. • If an NS, RS, RA, or Redirect message with a Link Layer Address (MAC) is received from a neighbor that is not yet in the ND cache, a new neighbor entry is created in the cache to store the received Link Layer MAC. The neighbor is put in the STALE state. This is the standard RFC behavior.
Parameters	global — Learns global neighbor entries out of received unsolicited Neighbor Advertisement messages. link-local — Learns link local neighbor entries out of received unsolicited Neighbor Advertisement messages. both — Learns both global and link local neighbor entries out of received unsolicited Neighbor Advertisement messages.

nd-populate-host-route

Syntax	[no] nd-populate-host-route
Context	config>service>vprn>if>ipv6
Description	This command enables the addition or deletion of host routes in the route-table derived from neighbor entries in the neighbor cache. To enable this command, the interface must be shut down. The command triggers the population of host routes in the route table out of their corresponding static, dynamic, or EVPN types in the neighbor table. Neighbor entries installed by subscriber management, local interfaces, and others, do not create host-routes. Only reachable entries are added to the route table (entries created from solicited NA messages). Entries created as STALE - from Neighbor Solicitation (NS), unsolicited Neighbor Advertisements (NA), Router Solicitation (RS), Router Advertisement (RA), and Redirect messages are not added to the route table since the neighbor is not confirmed as two-way. • RA, RS, NS, and Redirect messages with a link layer address are added as STALE cache entries. Unsolicited NAs are added as STALE if nd-learn-unsolicited is configured. • To speed up the addition of host routes to the route table for neighbors created as STALE, the following procedure is used: – If nd-populate-host-route is configured, the router sends an NS (unicast Neighbor Unreachability Detection (NUD) message) to the neighbor created as STALE. Only one NUD message is sent.

- If **nd-populate-host-route** is not configured, no confirmation message is sent and regular procedures apply.
- Only when the solicited NA for the neighbor is received the entry goes to REACHABLE and is then added to the route-table.

The **no** form of this command disables the creation of host routes from the neighbor cache.

nd-proactive-refresh

Syntax	nd-proactive-refresh { global link-local both } no nd-proactive-refresh
Context	config>service>vprn>if>ipv6
Description	This command enables a proactive refresh of the neighbor entries. When enabled, at the stale timer expiration, the router sends an NUD message to the host (regardless of the existence of traffic to the IP address on the IOM), so the entry can be refreshed or removed. This behavior is different from ARP, where the refresh is sent 30 seconds prior to the entry's age out time. The refresh can be optionally enabled for global addresses, link-local addresses, or both. The no form of the command disables the proactive behavior and the router only refreshes an entry if there is traffic that needs to be sent to the IP address.
Parameters	global — Refreshes global neighbor entries. This parameter is relevant only to global IPv6 addresses. link-local — Refreshes link local neighbor entries. This parameter is relevant only to global IPv6 addresses. both — Refreshes both global and link local neighbor entries. This parameter is relevant only to global IPv6 addresses.

nd-route-tag

Syntax	nd-route-tag <i>tag</i> no nd-route-tag
Context	config>service>vprn>if>ipv6
Description	This command adds a route tag to the ARP-ND host routes generated out of the neighbor entries in the interface. As any other route tag, it can be used to match ARP-ND routes in BGP export policies. The no form of this command removes the route tag for the ARP-ND host routes.

Parameters	<i>tag</i> — Specifies the route tag to be added when the proxy ND entries are advertised to EVPN.
Values	1 to 255

neighbor

Syntax	neighbor <i>ipv6-address mac-address</i> no neighbor <i>ipv6-address</i>
Context	config>service>vprn>if>ipv6
Description	This command configures IPv6-to-MAC address mapping on the interface.
Parameters	<i>ipv6-address</i> — Specifies the IPv6 address on the interface.
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
	<i>mac-address</i> — Specifies the 48-bit MAC address for the static ARP in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> where <i>aa</i> , <i>bb</i> , <i>cc</i> , <i>dd</i> , <i>ee</i> and <i>ff</i> are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

neighbor-limit

Syntax	neighbor-limit <i>limit</i> [log-only] [threshold percent] no neighbor-limit
Context	config>service>vprn>if>ipv6
Description	This command configures the maximum amount of dynamic IPv6 neighbor entries that can be learned on an IP interface. When the number of dynamic neighbor entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations will be dropped. Entries that have already been learned will be refreshed. The no form of the command removes the neighbor-limit.
Default	90 percent

Parameters	log-only — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit will be learned. <i>percent</i> — The threshold value (as a percentage) that triggers a warning message to be sent. Values 0 to 100 <i>limit</i> — The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic neighbor learning is disabled and no dynamic neighbor entries are learned. Values 0 to 102400

proxy-nd-policy

Syntax	proxy-nd-policy <i>policy-name</i> [<i>policy-name...</i> (up to 5 max)] no proxy-nd-policy
Context	config>service>vprn>if>ipv6
Description	This command configures a proxy neighbor discovery policy for the interface.
Parameters	<i>policy-name</i> — Specifies the existing policy name(s).

load-balancing

Syntax	load-balancing
Context	config>service>vprn>if config>service>vprn>nw-if
Description	This command enables the load-balancing context to configure interface per-flow load balancing options that will apply to traffic entering this interface and egressing over a LAG/ECMP on system-egress. This is a per interface setting. For load-balancing options that can also be enabled on the system level, the options enabled on the interface level overwrite system level configurations.
Default	not applicable

egr-ip-load-balancing

Syntax	egr-ip-load-balancing { <i>source</i> <i>destination</i> <i>inner-ip</i> } no egr-ip-load-balancing
Context	config>service>vprn>if>load-balancing config>service>vprn>if>nw-if>load-balancing

Description	This command specifies whether to include the source address or destination address or both in the LAG/ECMP hash on IP interfaces. Additionally, when I4-load-balancing is enabled, the command also applies to the inclusion of source/destination port in the hash inputs. The no form of this command includes both source and destination parameters.
Default	no egr-ip-load-balancing
Parameters	<i>source</i> — Specifies using the source address and (if I4-load balancing is enabled) source port in the hash, ignore destination address/port. <i>destination</i> — Specifies using the destination address and (if I4-load balancing is enabled) destination port in the hash, ignore source address/port. <i>inner-ip</i> — Specifies use of the inner IP header parameters instead of outer IP header parameters in LAG/ECMP hash for IPv4 encapsulated traffic.

lsr-load-balancing

Syntax	lsr-load-balancing <i>hashing-algorithm</i> no lsr-load-balancing
Context	config>service>vprn>nw-if>load-balancing
Description	This command specifies whether the IP header is used in the LAG and ECMP LSR hashing algorithm. This is the per interface setting.
Default	no lsr-load-balancing
Parameters	lbl-only — Only the label is used in the hashing algorithm. lbl-ip — The IP header is included in the hashing algorithm. ip-only — The IP header is used exclusively in the hashing algorithm. eth-encap-ip — The hash algorithm parses down the label stack (up to 3 labels supported) and once it hits the bottom, the stack assumes Ethernet II non-tagged header follows. At the expected Ethertype offset location, algorithm checks whether the value present is IPv4/v6 (0x0800 or 0x86DD). If the check passes, the hash algorithm checks the first nibble at the expected IP header location for IPv4/IPv6 (0x0100/0x0110). If the secondary check passes, the hash is performed using IP SA/DA fields in the expected IP header; otherwise (if any of the checks failed) label-stack hash is performed.

spi-load-balancing

Syntax	[no] spi-load-balancing
Context	config>service>vprn>if>load-balancing config>service>vprn>nw-if>load-balancing

Description	This command enables use of the SPI in hashing for ESP/AH encrypted IPv4/v6 traffic. This is a per interface setting. The no form disables the SPI function.
Default	no spi-load-balancing

teid-load-balancing

Syntax	[no] teid-load-balancing
Context	config>service>vprn>if>load-balancing config>service>vprn>nw-if>load-balancing
Description	This command enables inclusion of TEID in hashing for GTP-U/C encapsulates traffic for GTPv1/GTPv2. The no form of this command ignores TEID in hashing.
Default	no teid-load-balancing

local-proxy-arp

Syntax	[no] local-proxy-arp
Context	config>service>vprn>if config>service>vprn>sub-if>grp-if config>service>vprn>nw-if
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 thus will become the forwarding point for all traffic between hosts in that subnet. When local-proxy-arp is enabled, ICMP redirects on the ports associated with the service are automatically blocked.
Default	no local-proxy-arp

loopback

Syntax	[no] loopback
Context	config>service>vprn>if config>service>vprn>nw-if
Description	This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated interface cannot be bound to a SAP.

When using mtrace/mstat in a Layer 3 VPN context then the configuration for the VPRN should have a loopback address configured which has the same address as the core instance's system address (BGP next-hop).

Default no loopback

mac

Syntax [no] **mac** *ieee-mac-address*

Context config>service>vprn>if
config>service>vprn>if>vrrp
config>service>vprn>sub-if>grp-if
config>service>vprn>nw-if

Description This command assigns a specific MAC address to a VPRN IP interface.

The **no** form of this command returns the MAC address of the IP interface to the default value.

Default The physical MAC address associated with the Ethernet interface that the SAP is configured on.

Parameters *ieee-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.

monitor-oper-group

Syntax **monitor-oper-group** *name*
no monitor-oper-group

Context config>service>vprn>if

Description This command specifies the operational group to be monitored by the object under which it is configured. The oper-group name must be already configured under the config>service context before its name is referenced in this command.

The **no** form of the command removes the association from the configuration.

Default no monitor-oper-group

Parameters *name* — Specifies a character string of maximum 32 ASCII characters identifying the group instance.

proxy-arp

Syntax	[no] proxy-arp
Context	config>service>vprn>nw-if
Description	This command enables proxy ARP on the interface.
Default	no proxy-arp

proxy-arp-policy

Syntax	[no] proxy-arp-policy <i>policy-name</i> [<i>policy-name</i>]
Context	config>service>vprn>if config>service>vprn>sub-if>grp-if config>service>vprn>nw-if
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-policy
Parameters	<i>policy-name</i> — The export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. The specified name(s) must already be defined. Up to 5 route policies can be specified.

ptp-hw-assist

Syntax	[no] ptp-hw-assist
Context	config>service>vprn>if
Description	This command configures the 1588 port based timestamping assist function for the interface. This capability is supported on a specific set of hardware. The command may be blocked if not all hardware has the required level of support. If the SAP configuration of the interface is removed, the ptp-hw-assist configuration will be removed. If the IPv4 address configuration of the interface is removed, the ptp-hw-assist configuration will be removed. Only one interface per physical port can have ptp-hw-assist enabled.
Default	no ptp-hw-assist

qos-route-lookup

Syntax	qos-route-lookup [source destination] no qos-route-lookup
Context	config>service>vprn>if config>service>vprn>if>ipv6 config>service>vprn>sub-if>group-interface config>service>vprn>sub-if>grp-if>ipv6
Description	This command enables QoS classification of the ingress IP packets on an interface based on the QoS information associated with routes in the forwarding table. If the optional destination parameter is specified and the destination address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the destination address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network QoS policy. If the optional source parameter is specified and the source address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the source address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network QoS policy. If neither the optional source or destination parameter is present, then the default is destination address matching. The functionality enabled by the qos-route-lookup command can be applied to IPv4 packets or IPv6 packets on an interface, depending on whether it is present at the interface context (applies to IPv4) or the interface>ipv6 context (applies to IPv6). The ability to specify source address based QoS lookup is not supported for IPv6. For the 7740 ESS, subscriber management group interfaces also do not support the source QPPB option. The no form of the command reverts to the default.
Default	destination
Parameters	source — Enables QoS classification of incoming IP packets based on the source address matching a route with QoS information. destination — Enables QoS classification of incoming IP packets based on the destination address matching a route with QoS information.

redundant-interface

Syntax **redundant-interface** *red-ip-int-name*

no redundant-interface

Context	config>service>vprn config>service>vprn>sub-if>grp-if
Description	This command configures a redundant interface used for dual homing.
Parameters	<i>red-ip-int-name</i> — Specifies the redundant IP interface name.

remote-proxy-arp

Syntax	[no] remote-proxy-arp
Context	config>service>vprn>if config>service>vprn>sub-if>grp-if config>service>vprn>nw-if
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.
Default	no remote-proxy-arp

secondary

Syntax	secondary { <i>ip-address/mask</i> <i>ip-address netmask</i> } [broadcast all-ones host-ones] [igp-inhibit] no secondary { <i>ip-address/mask</i> <i>ip-address netmask</i> }
Context	config>service>vprn>if config>service>vprn>nw-if
Description	This command assigns a secondary IP address to the interface. Up to 16 total primary and secondary IPv4 and IPv6 addresses can be assigned to network interfaces, and up to 256 to access interfaces. Each address can be configured in an IP address, IP subnet or broadcast address format.



Caution: Configurations must not exceed 16 secondary IP addresses when IPsec, GRE, L2TPv3, or IP in IP protocols are active on an access interface.

-
- Parameters**
- ip-address* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).
 - mask* — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical 'AND' function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. A mask of 255.255.255.255 is reserved for system IP addresses.
 - netmask* — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.
 - broadcast** — The optional **broadcast** parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**. The broadcast format on an IP interface can be specified when the IP address is assigned or changed. This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) will be received by the IP interface. (Default: *host-ones*)
 - all-ones** — The **all-ones** keyword following the **broadcast** parameter specifies the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.
 - host-ones** — The **host-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default used by an IP interface.
The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.
 - igp-inhibit** — The optional **igp-inhibit** parameter signals that the given secondary IP interface should not be recognized as a local interface by the running IGP. For OSPF and IS-IS, this means that the specified secondary IP interfaces will not be injected and used as passive interfaces and will not be advertised as internal IP interfaces into the IGP's link state database. For RIP, this means that these secondary IP interfaces will not source RIP updates.

shcv-policy-ipv4

Syntax	shcv-policy-ipv4 <i>policy-name</i> no shcv-policy-ipv4
Context	config>service>vprn>if config>service>vprn>sub-if>grp-if
Description	This command specifies the Subscriber Host Connectivity Verification (SHCV) policy for IPv4 only. The no form of the command removes the policy name from the SAP configuration.

shcv-policy-ipv6

Syntax	shcv-policy-ipv6 <i>policy-name</i> no shcv-policy-ipv6
Context	config>service>vprn>if config>service>vprn>sub-if>grp-if
Description	This command specifies the Subscriber Host Connectivity Verification (SHCV) policy for IPv6 only. The no form of the command removes the policy name from the SAP configuration.

static-arp

Syntax	static-arp <i>ieee-mac-address unnumbered</i> no static-arp <i>unnumbered</i>
Context	config>service>vprn>if config>service>vprn>nw-if
Description	This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP will appear in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface. If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address. The no form of this command removes a static ARP entry.
Parameters	<i>ip-address</i> — Specifies the IP address for the static ARP in IP address dotted decimal notation.

ieee-mac-address — Specifies the 48-bit MAC address for the static ARP in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

unnumbered — Specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. Once this command is configured, it overrides any dynamic ARP.

static-tunnel-redundant-next-hop

Syntax	static-tunnel-redundant-next-hop <i>ip-address</i> no static-tunnel-redundant-next-hop
Context	config>service>vprn>if
Description	This command specifies redundant next-hop address on public or private IPsec interface (with public or private tunnel-sap) for static IPsec tunnel. The specified next-hop address will be used by standby node to shunt traffic to master in case of it receives them. The next-hop address will be resolved in routing table of corresponding service. The no form of the command removes the address from the interface configuration.
Parameters	<i>ip-address</i> — Specifies the static ISA tunnel redundant next-hop address.

secure-nd

Syntax	[no] secure-nd
Context	config>service>vprn>if>ipv6
Description	This command enables Secure Neighbor Discovery (SeND) on the IPv6 interface. The no form of the command reverts to the default and disabled SeND.

allow-unsecured-msgs

Syntax	[no] allow-unsecured-msgs
Context	config>service>vprn>if>secure-nd
Description	This command specifies whether unsecured messages are accepted. When Secure Neighbor Discovery (SeND) is enabled, only secure messages are accepted by default. The no form of the command disables accepting unsecured messages.

link-local-modifier

Syntax	link-local-modifier <i>modifier</i> [no] link-local-modifier
Context	config>service>vprn>if>secure-nd
Description	This command configures the Cryptographically Generated Address (CGA) modifier for link-local addresses.
Parameters	<i>modifier</i> — Specifies the modifier in 32 hexadecimal nibbles. Values 0x0–0xFFFFFFFF

public-key-min-bits

Syntax	public-key-min-bits <i>bits</i> [no] public-key-min-bits
Context	config>service>vprn>if>secure-nd
Description	This command configures the minimum acceptable key length for public keys used in the generation of a Cryptographically Generated Address (CGA).
Parameters	<i>bits</i> — Specifies the number of bits. Values 512–1024

security-parameter

Syntax	security-parameter <i>sec</i> [no] security-parameter
Context	config>service>vprn>if>secure-nd
Description	This command configures the security parameter used in the generation of a Cryptographically Generated Address (CGA).
Parameters	<i>sec</i> — Specifies the security parameter. Values 0 to 1

shutdown

Syntax	[no] shutdown
Context	config>service>vprn>if>secure-nd

Description This command enables or disables Secure Neighbor Discovery (SeND) on the interface.

stale-time

Syntax **stale-time** *seconds*
no stale-time

Context config>service>vprn>ipv6
config>service>vprn>if>ipv6

Description This command configures the time a neighbor discovery cache entry can remain stale before being removed.

The **no** form of the command removes the stale-time value.

Default no stale-time

Parameters *seconds* — The allowed stale time (in seconds) before a neighbor discovery cache entry is removed.

Values 60 to 65535

tcp-mss

Syntax **tcp-mss** *mss-value*
no tcp-mss

Context service>vprn>if
service>vprn>if>ipv6

Description This command statically sets the TCP maximum segment size (MSS) for TCP connections originated from the associated IP interface to the specified value.

The **no** form of the command removes the static value and allows the TCP MSS value to be calculated based on the IP MTU value by subtracting the base IP and TCP header lengths from the IP MTU value ($\text{tcp\_mss} = \text{ip\_mtu} - 40$).

Default no tcp-mss

Parameters *mss-value* — The TCP MSS value that should be used in the TCP SYN packet during the three-way handshake negotiation of a TCP connection.

Note: $9158 = \text{max-IP\_MTU} (9198) - 40$

Values 536 to 9158 (IPv4)
1220 to 9138 (IPv6)

tos-marking-state

Syntax	tos-marking-state {trusted untrusted} no tos-marking-state
Context	config>service>vprn>sub-if>grp-if config>service>vprn>nw-if
Description	This command is used to alter the default trusted state to a non-trusted state. When unset or reverted to the trusted default, the ToS field will not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set, in which case the egress network interface treats all VPRN and network IP interface as untrusted. When the ingress interface is set to untrusted, all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface. The egress network remarking rules also apply to the ToS field of IP packets routed using IGP shortcuts (tunneled to a remote next-hop). However, the tunnel QoS markings are always derived from the egress network QoS definitions. Egress marking and remarking is based on the internal forwarding class and profile state of the packet once it reaches the egress interface. The forwarding class is derived from ingress classification functions. The profile of a packet is either derived from ingress classification or ingress policing. The default marking state for network IP interfaces is trusted. This is equivalent to declaring no tos-marking-state on the network IP interface. When undefined or set to tos-marking-state trusted, the trusted state of the interface will not be displayed when using show config or show info unless the detail parameter is given. The save config command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified. The no tos-marking-state command is used to restore the trusted state to a network IP interface. This is equivalent to executing the tos-marking-state trusted command.
Default	tos-marking-state trusted
Parameters	trusted — The default prevents the ToS field to not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set. untrusted — Specifies that all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface.

ipv6

Syntax	[no] ipv6
Context	config>service>vprn>sub-if config>service>vprn>sub-if>grp-if
Description	This command configures IPv6 parameters.

allow-unmatching-prefixes

Syntax	[no] allow-unmatching-prefixes
Context	config>service>vprn>sub-if
Description	This command allows address assignment to PPPoX hosts in cases where the assigned address falls outside the range of the configured subnets below the subscriber interface. Alternatively, if the interface is configured as unnumbered, this command cannot be enabled.
Default	no allow-unmatching-prefixes

allow-unmatching-prefixes

Syntax	[no] allow-unmatching-prefixes
Context	config>service>vprn>sub-if>ipv6> config>service>ies>sub-if>ipv6>
Description	This command allows address assignment to IPv6 hosts in cases where the assigned address or the prefix falls outside of the range of the configured IPv6 subscriber-prefixes under the config>service>vprn/ies>sub-if>ipv6 hierarchy. Unnumbered PPPoEv6 does not mean that the PPPoEv6 hosts do not have an IPv6 address or prefix assigned. It only means that the IPv6 address range (out of which the address or prefix is assigned to the host) does not have to be known in advance via configuration under the subscriber-interface>ipv6>subscriber-prefixes node.
Default	no allow-unmatching-prefixes

delegated-prefix-length

Syntax	delegated-prefix-length <i>bits</i> delegated-prefix-length variable no delegated-prefix-length
Context	config>router>sub-if>ipv6 config>service>vprn>sub-if>ipv6
Description	This command configures the subscriber-interface level setting for delegated prefix length. The delegated prefix length for a subscriber- interface can be either set to a fixed value that is explicitly configured under the subscriber-interface CLI hierarchy or a variable value that can be obtained from various sources. This command can be changed only when no IPv6 prefixes are configured under the subscriber-interface.
Default	no delegated-prefix-length This means that the delegated prefix length is 64.

Parameters *bits* — The delegated prefix length in bits. This value will be applicable to the entire subscriber-interface. In case that the delegated prefix length is also supplied via other means (LUDB, Radius or DHCP Server), such supplied value must match the value configured under the subscriber-interface. Otherwise the prefix instantiation in the router fails.

Values 48 to 64

variable — The delegated prefix value can be of any length between 48 to 64. The value itself can vary between the prefixes and it will be provided at the time of prefix instantiation. The order of priority for the source of the delegated prefix length is:

- LUDB
- Radius
- DHCPv6 server

dhcp6

Syntax [no] dhcp6

Context config>service>vprn>sub-if>grp-if>ipv6

Description This command allows access to the DHCP6 context within the group interface configuration. Within this context, DHCP6 parameters can be configured.

Default no dhcp6

option

Syntax [no] option

Context config>service>vprn>sub-if>grp-if>ipv6

Description This command enters the context to configure DHCPv6 relay information options.

The **no** form of the command disables DHCPv6 relay information options.

interface-id

Syntax interface-id
interface-id ascii-tuple
interface-id ifindex
interface-id sap-id
interface-id string
no interface-id

Context config>service>vprn>sub-if>grp-if>ipv6>dhcp6>option

Description	This command enables the sending of interface ID options in the DHCPv6 relay packet. The no form of the command disables the sending of interface ID options in the DHCPv6 relay packet
Parameters	ascii-tuple — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “ ”. ifindex — Specifies that the interface index will be used (the If Index of a router interface can be displayed using the command show>router>if>detail). sap-id — Specifies that the SAP identifier will be used. string — Specifies a string of up to 32 characters long, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

remote-id

Syntax	[no] remote-id
Context	config>service>vprn>sub-if>grp-if>ipv6>dhcp6>option
Description	This command enables the sending of remote ID option in the DHCPv6 relay packet. The client DHCP Unique Identifier (DUID) is used as the remote ID. The no form of the command disables the sending of remote ID option in the DHCPv6 relay packet.

proxy-server

Syntax	[no] proxy-server
Context	config>service>vprn>sub-if>grp-if>ipv6>dhcp6
Description	This command allows access to the DHCP6 proxy server context. Within this context, DHCP6 proxy server parameters of the group interface can be configured.
Default	no proxy-server

renew-timer

Syntax	renew-timer [days days] [hrs hours] [min minutes] [sec seconds] no renew-timer
Context	config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server

Description This command configures the renew-timer (T1), the time at which the client contacts the server from which the addresses in the IA_NA or IA_PD were obtained to extend the lifetimes of the addresses or prefixes assigned to the client.

Default renew-timer min 30

Parameters *time* — Specifies the renew timer.

Values

days:	0 to 7
hours	0 to 23
minutes:	0 to 59
seconds	0 to 59

rebind-timer

Syntax **rebind-timer** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no rebind-timer

Context config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server

Description This command configures the rebind-timer (T2), the time at which the client contacts any available server to extend the lifetimes of the addresses or prefixes assigned to the client.

Default rebind-timer min 48

Parameters *time* — Specifies the rebind timer.

Values

days:	0 to 14
hours	0 to 23
minutes:	0 to 59
seconds	0 to 59

preferred-lifetime

Syntax **preferred-lifetime infinite**
preferred-lifetime [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no preferred-lifetime

Context config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server

Description The preferred lifetime for the IPv6 prefix or address in the option, expressed in units of seconds. When the preferred lifetime expires, any derived addresses are deprecated.

Default preferred-lifetime hrs 1

Parameters **infinite** — Sets the preferred lifetime to infinite value.

time — Specifies the preferred lifetime timer.

Values

days:	0 to 49710
hours	0 to 23
minutes:	0 to 59
seconds	0 to 59

valid-lifetime

Syntax **valid-lifetime infinite**

valid-lifetime [*days days*] [*hrs hours*] [*min minutes*] [*sec seconds*]

no valid-lifetime

Context config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server

Description The valid lifetime for the IPv6 prefix or address in the option, expressed in units of seconds.

Default valid-lifetime days 1

Parameters **infinite** — Sets the valid lifetime to infinite value.

time — Specifies the valid lifetime timer.

Values

days:	0 to 49710
hours	0 to 23
minutes:	0 to 59
seconds	0 to 59

client-applications

Syntax **client-applications** [*dhcp*] [*ppp*]

no client-applications

Context config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server

Description This command configures the client host types to which the DHCP6 proxy server is allowed to assign addresses.

Parameters **dhcp** — Specifies IP over Ethernet hosts.

ppp — Specifies PPP over Ethernet hosts.

router-advertisements

Syntax	[no] router-advertisements
Context	config>service>vprn>sub-if>grp-if>ipv6
Description	This command enables Router Advertisement transmission on this group interface.
Default	router-advertisements

current-hop-limit

Syntax	current-hop-limit <i>hop-count</i> no current-hop-limit
Context	config>service>vprn>sub-if>grp-if>ipv6>router-ad
Description	This command specifies the hop-limit advertised to hosts in router advertisements.
Default	current-hop-limit 64
Parameters	<i>hop-count</i> — Specifies the current hop limit (decimal) inserted into router advertisements. Values 0 to 255

managed-configuration

Syntax	[no] managed-configuration
Context	config>service>vprn>sub-if>grp-if>ipv6>router-ad
Description	This command sets the managed address configuration flag. This flag indicates that DHCPv6 is available for address configuration in addition to any address auto-configured using stateless address auto-configuration. See RFC 3315, Dynamic Host Configuration Protocol (DHCP) for IPv6.
Default	no managed-configuration

max-advertisement

Syntax	max-advertisement <i>seconds</i> no max-advertisement
Context	config>service>vprn>sub-if>grp-if>ipv6>rtr-adv config>service>vprn>sub-if>ipv6>rtr-adv

Description	This command configures the maximum interval between sending router advertisement messages.
Default	max-advertisement 1800
Parameters	<i>seconds</i> — Specifies the maximum interval in seconds between sending router advertisement messages.
Values	900 to 1800

min-advertisement

Syntax	min-advertisement <i>seconds</i> no min-advertisement
Context	config>service>vprn>sub-if>grp-if>ipv6>rtr-adv config>service>vprn>sub-if>ipv6>rtr-adv
Description	This command configures the minimum interval between sending router advertisement messages.
Default	min-advertisement 900
Parameters	<i>seconds</i> — Specifies the minimum interval, in seconds, between sending router advertisement messages.
Values	900 to 1350

mtu

Syntax	mtu <i>bytes</i> no mtu
Context	config>service>vprn>sub-if>grp-if>ipv6>router-ad
Description	This command configures the MTU for the nodes to use to send packets on the link.
Default	no mtu
Parameters	<i>bytes</i> — Specifies the MTU for the nodes to use to send packets on the link.
Values	1280 to 9212

other-stateful-configuration

Syntax	[no] other-stateful-configuration
Context	config>service>vprn>sub-if>grp-if>ipv6>router-ad

Description	This command sets the “Other configuration” flag. This flag indicates that DHCPv6 is available for auto-configuration of other (non-address) information such as DNS-related information or information on other servers in the network. See RFC 3736, Stateless Dynamic Host Configuration Protocol (DHCP) for IPv6.
Default	no other-stateful-configuration

prefix-options

Syntax	[no] prefix-options
Context	config>service>vprn>sub-if>grp-if>ipv6>router-ad
Description	This command configures router advertisement parameters for IPv6 prefixes returned via RADIUS Framed-IPv6-Prefix. All prefixes will inherit these configuration parameters.
Default	no prefix-options

autonomous

Syntax	[no] autonomous
Context	config>services>vprn>sub-if>grp-if>ipv6>router-ad>prefix-op
Description	This command specifies whether the prefix can be used for stateless address auto-configuration.
Default	no autonomous

preferred-lifetime

Syntax	preferred-lifetime [<i>seconds</i> <i>infinite</i>] no preferred-lifetime
Context	config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>prefix-options
Description	This command configures the remaining length of time in seconds that this prefix will continue to be preferred, such as, time until deprecation. The address generated from a deprecated prefix should not be used as a source address in new communications, but packets received on such an interface are processed as expected.
Default	preferred-lifetime 3600
Parameters	<i>seconds</i> — Specifies a decimal time interval in seconds. Values 0 to 4294967295 <i>infinite</i> — Specifies a 0xffffffff value. Dec = 4294967295.

valid-lifetime

Syntax	valid-lifetime [<i>seconds</i> infinite] no valid-lifetime
Context	config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>prefix-options
Description	This command specifies the length of time, in seconds, that the prefix is valid for the purpose of on-link determination. A value of all one bits (0xffffffff) represents infinity. The address generated from an invalidated prefix should not appear as the destination or source address of a packet.
Default	valid-lifetime 86400
Parameters	<i>seconds</i> — Specifies a decimal time interval in seconds. Values 0 to 4294967295 infinite — Specifies a 0xffffffff value. Dec = 4294967295.

reachable-time

Syntax	reachable-time <i>milliseconds</i> no reachable-time
Context	config>services>vprn>sub-if>grp-if>ipv6>rtr-adv
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	<i>milliseconds</i> — The length of time the router should be considered reachable for default router selection. Values 0-3600000

retransmit-time

Syntax	retransmit-time <i>milliseconds</i> no retransmit-time
Context	config>services>vprn>sub-if>grp-if>ipv6>rtr-adv
Description	This command configures the retransmission frequency of neighbor solicitation messages.
Default	no retransmit-time

Parameters	<i>milliseconds</i> — Specifies how often the retransmission should occur.
Values	0 to 1800000

router-lifetime

Syntax	router-lifetime <i>seconds</i> router-lifetime no-default-router no router-lifetime
Context	config>services>vprn>sub-if>grp-if>ipv6>rtr-adv
Description	This command sets the router lifetime. A value of zero indicates this router should not be used by hosts as a default router.
Default	router-lifetime 4500
Parameters	<i>seconds</i> — Specifies how long this router is valid for default router selection.
Values	2700 to 9000

delegated-prefix-length

Syntax	[no] delegated-prefix-length <i>prefix-length</i>
Context	config>service>vprn>sub-if>ipv6
Description	This command defines the prefix-length used for all DHCPv6 prefix delegations on this subscriber interface.
Parameters	<i>prefix-length</i> — Specifies the prefix length in use on this subscriber interface for DHCPv6 IA_PD.
Values	48 to 64
Default	64

subscriber-prefixes

Syntax	subscriber-prefixes
Context	config>service>vprn>sub-if>ipv6
Description	This command specifies aggregate off-link subscriber prefixes associated with this subscriber interface. Individual prefixes are specified under the prefix context list aggregate routes in which the next-hop is indirect via the subscriber interface.

prefix

Syntax	prefix <i>ipv6-address/prefix-length</i> [pd] [wan-host] no prefix <i>ipv6-address/prefix-length</i>
Context	config>service>vprn>sub-if>ipv6>sub-prefixes
Description	This command allows a list of prefixes (using the prefix command multiple times) to be routed to hosts associated with this subscriber interface. Each prefix will be represented in the associated FIB with a reference to the subscriber interface. Prefixes are defined as being for prefix delegation (pd) or use on a WAN interface or host (wan-host).
Parameters	<i>ipv6-address</i> — Specifies the 128-bit IPv6 address. Values 128-bit hexadecimal IPv6 address in compressed form <i>prefix-length</i> — Specifies the length of any associated aggregate prefix. Values 32 to 63 pd — Specifies that this aggregate is used by IPv6 ESM hosts for DHCPv6 prefix-delegation. wan-host — Specifies that this aggregate is used by IPv6 ESM hosts for local addressing or by a routing gateway's WAN interface.

private-retail-subnets

Syntax	[no] private-retail-subnets
Context	config>service>vprn>sub-if
Description	This command controls the export of subnets to the forwarding service. When this attribute is configured, subnets defined on this retail subscriber interface will no longer be exported to the associated wholesale VPRN and will remain private to the retail VPRN. This is useful in a PPPoE business service context as it allows retail services to use overlapping IP address spaces even if these services are associated with the same wholesale service. PPPoE sessions are actually terminated in the retail service although their traffic transits on a SAP belonging to the wholesale service. This configuration is incompatible, however, with IPoE host management (DHCP, static-host and ARP-host) as these host types require that the retail subnets are exported to the wholesale VPRN. Thus, if PPPoE sessions need to coexist with IPoE hosts, this attribute should not be configured on this retail interface. This command fails if the subscriber interface is not associated with a wholesale service. If the retail VPRN is of the type hub, this attribute is mandatory. Then, it will be enabled by default and it will not be possible to deconfigure it.

unnumbered

Syntax	unnumbered [<i>ip-int-name</i> <i>ip-address</i>] no unnumbered
Context	config>service>vprn>if config>service>vprn>nw-if
Description	This command configures the interface as an unnumbered interface. An unnumbered IP interface is supported on a SONET/SDH access port with the PPP, ATM, Frame Relay, cisco-HDLC encapsulation. It is not supported on access ports that do not carry IP traffic, but are used for native TDM circuit emulation.
Parameters	<i>ip-int-name</i> — Specifies the name of an IP interface. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. <i>ip-address</i> — Specifies an IP address.

qos

Syntax	qos <i>network-policy-id</i> port-redirect-group <i>queue-group-name</i> egress-instance <i>instance-id</i> fp-redirect-group <i>queue-group-name</i> ingress-instance <i>instance-id</i> no qos
Context	config>service>vprn>nw-if
Description	This command associates a network Quality of Service (QoS) policy with a network IP interface. Only one network QoS policy can be associated with an IP interface at one time. Attempts to associate a second QoS policy return an error. Associating a network QoS policy with a network interface is useful for the following purposes: • To apply classification rules for determining the forwarding-class and profile of ingress packets on the interface. • To associate ingress packets on the interface with a queue-group instance applied to the ingress context of the interface's forwarding plane (FP). (This is only applicable to interfaces on IOM3 and later cards.) The referenced ingress queue-group instance may have policers defined in order to rate limit ingress traffic on a per-forwarding class (and forwarding type: unicast vs. multicast) basis. • To perform 802.1p, DSCP, IP precedence and/or MPLS EXP re-marking of egress packets on the interface. • To associate egress packets on the interface with a queue-group instance applied to the egress context of the interface's port. The referenced egress queue-group instance may have policers and/or queues defined in order to rate limit egress traffic on a per-forwarding class basis.

The **no** form of the command removes the network QoS policy association from the network IP interface, and the QoS policy reverts to the default.

Default	no qos
Parameters	<i>network-policy-id</i> — An existing network policy ID to associate with the IP interface.
Values	1 to 65535
port-redirect-group	<i>queue-group-name</i> — This optional parameter specifies the egress queue-group used for all egress forwarding-class redirections specified within the network QoS policy ID. The specified <i>queue-group-name</i> must exist as an egress queue group applied to the egress context of the port associated with the IP interface.
egress-instance	<i>instance-id</i> — Since multiple instances of the same egress queue-group can be applied to the same port this optional parameter is used to specify which particular instance to associate with this particular network IP interface.
Values	1 to 16384
fp-redirect-group	<i>queue-group-name</i> — This optional parameter specifies the ingress queue-group used for all ingress forwarding-class redirections specified within the network QoS policy ID. The specified <i>queue-group-name</i> must exist as an ingress queue group applied to the ingress context of the forwarding plane associated with the IP interface.
ingress-instance	<i>instance-id</i> — Since multiple instances of the same ingress queue-group can be applied to the same forwarding plane this parameter is required to specify which particular instance to associate with this particular network IP interface.
Values	1 to 16384

urpf-check

Syntax	[no] urpf-check
Context	config>service>vprn>if config>service>vprn>nw-if config>service>vprn>if>ipv6 config>service>vprn>sub-if>grp-if
Description	This command enables unicast RPF (uRPF) check on this interface. The no form of the command disables unicast RPF (uRPF) Check on this interface.
Default	no urpf-check

vas-if-type

Syntax	vas-if-type {to-from-access to-from-network to-from-both} no vas-if-type

Context	config>service>vprn>if
Description	This command configures the type of a Value Added Service (VAS) facing interface. To change the vas-if-type, the shutdown command is required. The vas-if-type and loopback commands are mutually exclusive. The no form of the command removes the VAS interface type configuration.
Default	no vas-if-type
Parameters	to-from-access — Used when two separate (to-from-access and to-from-network) interfaces are used for VAS connectivity. For service chaining, traffic arriving from access interfaces (upstream) is redirected to a PBR target reachable over this interface for upstream VAS processing. Downstream traffic after VAS processing must arrive on this interface, so that the traffic is subject to regular routing but is not subject to AA divert, nor egress subscriber PBR. to-from-network — Used when two separate (to-from-access and to-from-network) interfaces are used for VAS connectivity. For service chaining, traffic arriving from network interfaces (downstream) is redirected to a PBR target reachable over this interface for downstream VAS processing. Upstream traffic after VAS processing must arrive on this interface, so that regular routing can be applied. to-from-both — Used when a single interface is used for VAS connectivity (no local-to-local traffic). For service chaining, both traffic arriving from access and from network is redirected to a PBR target reachable over this interface for upstream/downstream VAS processing. Traffic after VAS processing must arrive on this interface, so that the traffic is subject to regular routing but is not subject to AA divert, nor egress subscriber PBR.

mode

Syntax	mode {strict loose strict-no-ecmp} no mode
Context	config>service>vprn>if>urpf-check config>service>vprn>nw-if>urpf-check config>service>vprn>sub-if>grp-if>urpf-check
Description	This command specifies the mode of unicast RPF check. The no form of the command reverts to the default (strict) mode.
Default	mode strict
Parameters	strict — When specified, uRPF checks whether incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix.

loose — In **loose** mode, uRPF checks whether incoming packet has source address with a corresponding prefix in the routing table. However, the loose mode does not check whether the interface expects to receive a packet with a specific source address prefix. This object is valid only when **urpf-check** is enabled.

strict-no-ecmp — When a packet is received on an interface in this mode and the SA matches an ECMP route the packet is dropped by uRPF.

3.8.2.17 Network Interface Commands

network-interface

Syntax	network-interface <i>interface-name</i> [create] no network-interface <i>interface-name</i>
Context	config>service>vprn
Description	This command configures a network interface in a VPRN that acts as a CSC interface to a CSC-CE in a Carrier Supporting Carrier IP VPN deployment model.

3.8.2.18 Interface DHCP Commands

dhcp

Syntax	dhcp
Context	config>service>vprn>if config>service>vprn>nw-if config>service>vprn>sub-if config>service>vprn>sub-if>grp-if
Description	This command enters the context to configure DHCP parameters.

client-applications

Syntax	client-applications dhcp ppp no client-applications
Context	config>service>vprn>sub-if>grp-if>dhcp config>service>vprn>sub-if>dhcp
Description	This command enables the clients that will try to contact the DHCP server(s).

The **no** form of the command removes the server client type from the configuration.

- Parameters** **dhcp** — Specifies that the DHCP relay will forward requests to the DHCP server(s).
ppp — Specifies that PPPoE will attempt to request an IP address for a PPPoE client from the DHCP server(s) assigned to PPPoE node.

action

- Syntax** **action {replace | drop | keep}**
no action
- Context** config>service>vprn>if>dhcp>option
 config>service>vprn>nw-if>dhcp>option
 config>service>vprn>sub-if>grp-if>dhcp>option
- 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 returns the system to the default value.
- Default** Per RFC 3046, *DHCP Relay Agent Information Option*, section 2.1.1, *Reforwarded DHCP requests*, the default is to keep the existing information intact. The exception to this is if the giaddr of the received packet is the same as the ingress address on the router. In that case the packet is dropped and an error is logged.
- Parameters** **replace** — In the upstream direction (from the user), the existing Option 82 field is replaced with the Option 82 field from the router. In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).
drop — The packet is dropped, and an error is logged.
keep — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is sent on towards the client.
- The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.
- If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.

circuit-id

- Syntax** **circuit-id [ascii-tuple | ifindex | sap-id | vlan-ascii-tuple]**
no circuit-id
- Context** config>service>vprn>if>dhcp>option

	<pre>config>service>vprn>nw-if>dhcp>option config>service>vprn>sub-if>grp-if>dhcp>option</pre>
Description	When enabled, the router sends the interface index (If Index) in the circuit-id suboption of the DHCP packet. The If Index of a router interface can be displayed using the command show>router>interface>detail. This option specifies data that must be unique to the router that is relaying the circuit. If disabled, the circuit-id suboption of the DHCP packet will be left empty. The no form of this command returns the system to the default.
Default	circuit-id ascii-tuple
Parameters	ascii-tuple — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “ ”. ifindex — Specifies that the interface index will be used. The If Index of a router interface can be displayed using the command show>router>interface>detail. sap-id — Specifies that the SAP ID will be used. vlan-ascii-tuple — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

filter

Syntax	<pre>filter filter-id no filter</pre>
Context	config>service>vprn>sub-if>grp-if>dhcp
Description	This command configures the DHCP filter for this interface.
Parameters	filter-id — Specifies the filter policy. The filter ID must already exist within the created IP filters. Values 1 to 65535

gi-address

Syntax	<pre>gi-address ip-address [src-ip-addr] no gi-address</pre>
Context	<pre>config>service>vprn>if>dhcp config>service>vprn>nw-if>dhcp config>service>vprn>sub-if>dhcp</pre>

```
config>service>vprn>sub-if>grp-if>dhcp
```

- 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.
- 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.

lease-populate

- Syntax** **lease-populate** [*nbr-of-leases*]
lease-populate [*nbr-of-leases*] **route-populate** [*pd*] *na* [*ta*]
lease-populate [*nbr-of-leases*] **route-populate** *pd* [*na*] [*ta*] [**exclude**]
lease-populate [*nbr-of-leases*] **route-populate** [*pd*] [*na*] *ta*
no lease-populate
- Context** config>service>vprn>if>ipv6>dhcp-relay
- Description** This command specifies the maximum number of DHCPv6 lease states allocated by the DHCPv6 relay function, allowed on this interface.
- Optionally, by specifying “route-populate” parameter, system could:
- Create routes based on the IA_PD/IA_NA/IA_TA prefix option in relay-reply message.
 - Create black hole routes based on OPTION_PD_EXCLUDE in IA_PD in relay-reply message.
- These routes could be redistributed into IGP/BGP by using route-policy, following protocol types that could be used in “from protocol”:
- dhcpv6-pd
 - dhcpv6-na
 - dhcpv6-ta
 - dhcpv6-pd-excl
- Parameters** *nbr-of-entries* — Defines the number lease state table entries allowed for this interface. If this parameter is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCPv6 ACK messages are discarded.
- Values** 1 to 8000

route-populate — Specifies the route populate on which to allocate DHCPv6 lease states.

Values pd/na/ta — Create route based on specified option.
 exclude — Create black hole route based on
 OPTION_PD_EXCLUDE.

neighbor-resolution

Syntax [no] neighbor-resolution

Context config>service>vprn>if>ipv6>dhcp6-relay

Description This command enables neighbor resolution with DHCPv6 relay.

The **no** form of the command disables neighbor resolution.

match-circuit-id

Syntax [no] match-circuit-id

Context config>service>vprn>sub-if>grp-if>dhcp

Description This command enables Option 82 circuit ID on relayed DHCP packet matching. For routed CO, the group interface DHCP relay process is stateful. When packets are relayed to the server the virtual router ID, transaction ID, SAP ID, and client hardware MAC address of the relayed packet are tracked.

When a response is received from the server the virtual router ID, transaction ID, and client hardware MAC address must be matched to determine the SAP on which to send the packet out. In some cases, the virtual router ID, transaction ID, and client hardware MAC address are not guaranteed to be unique.

When the **match-circuit-id** command is enabled we use this as part of the key to guarantee correctness in our lookup. This is really only needed when we are dealing with an IP aware DSLAM that proxies the client hardware MAC address.

Default no match-circuit-id

option

Syntax [no] option

Context config>service>vprn>if>dhcp
 config>service>vprn>nw-if>dhcp
 config>service>vprn>sub-if>dhcp
 config>service>vprn>sub-if>grp-if>dhcp

Description This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.

The **no** form of this command returns the system to the default.

Default no option

copy-82

Syntax [no] copy-82

Context config>service>vprn>nw-if>dhcp>option

Description This command enables the copy-82 option.

The **no** form of the command disables the option.

remote-id

Syntax remote-id [mac | string *string*]
no remote-id

Context config>service>vprn>sub-if>grp-if>dhcp>option
config>service>vprn>nw-if>dhcp>option

Description 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 returns the system to the default.

Default no remote-id

Parameters **mac** — 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>vprn>if>dhcp>option
config>service>vprn>nw-if>dhcp>option
config>service>vprn>sub-if>grp-if>dhcp>option

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>vprn>if>dhcp>option config>service>vprn>nw-if>dhcp>option config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor
Description	This command enables the sending of the MAC address in the Nokia vendor specific suboption of the DHCP relay packet. The no form of the command disables the sending of the MAC address in the Nokia vendor specific suboption of the DHCP relay packet.

pool-name

Syntax	[no] pool-name
Context	config>service>vprn>if>dhcp>option
Description	This command enables the sending of the pool name in the Nokia vendor-specific suboption of the DHCP relay packet. The no form of the command disables the feature.

if-name

Syntax	[no] if-name
Context	config>service>vprn>nw-if>dhcp>option
Description	This command enables the sending of the interface name in the Nokia vendor-specific suboption of the DHCP relay packet The no form of the command disables the sending.

port-id

Syntax	[no] port-id
Context	config>service>vprn>nw-if>dhcp>option
Description	This command enables sending of the port-id in the Nokia vendor-specific suboption of the DHCP relay packet The no form of the command disables the sending.

sap-id

Syntax	[no] sap-id
Context	config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor
Description	This command enables the sending of the SAP ID in the Nokia vendor-specific suboption of the DHCP relay packet. The no form of the command disables the sending of the SAP ID in the Nokia vendor specific suboption of the DHCP relay packet.

service-id

Syntax	[no] service-id
Context	config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor
Description	This command enables the sending of the service ID in the Nokia vendor specific suboption of the DHCP relay packet. The no form of the command disables the sending of the service ID in the Nokia vendor specific suboption of the DHCP relay packet.

string

Syntax	[no] string <i>text</i>
Context	config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor
Description	This command specifies the vendor specific suboption string of the DHCP relay packet. The no form of the command returns the default value.
Parameters	<i>text</i> — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (" ").

system-id

Syntax	[no] system-id
Context	config>service>vprn>if>dhcp>option>vendor config>service>vprn>nw-if>dhcp>option>vendor

config>service>vprn>sub-if>grp-if>dhcp>option>vendor

Description This command specifies whether the system-id is encoded in the Nokia vendor specific sub-option of Option 82.

proxy-server

Syntax **proxy-server**

Context config>service>if>dhcp
config>service>vprn>sub-if>grp-if>dhcp

Description This command configures the DHCP proxy server.

emulated-server

Syntax **emulated-server** *ip-address*
no emulated-server

Context config>service>vprn>if>dhcp>proxy
config>service>vprn>sub-if>grp-if>dhcp>proxy-server

Description This command configures the IP address to be used as the DHCP server address in the context of this service. Typically, the configured address should be in the context of the subnet.

The **no** form of this command reverts to the default setting. The local proxy server will not become operational without a specified emulated server address.

Parameters *ip-address* — Specifies the emulated server address.

Default For a retail interface, the default is the local interface.

lease-time

Syntax **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*] [**override**]
no lease-time

Context config>service>vprn>if>dhcp>proxy
config>service>vprn>sub-if>grp-if>dhcp>proxy-server

Description This command defines the length of lease-time that will be provided to DHCP clients. By default the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.

The no form of this command disables the use of the lease-time command. The local-proxy-server will use the lease-time offered by either a RADIUS or DHCP server.

Default	7 days 0 hours 0 seconds
Parameters	override — Specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients. days — Specifies the number of days that the given IP address is valid. Values 0 to 3650 hours — Specifies the number of hours that the given IP address is valid. Values 0 to 23 minutes — Specifies the number of minutes that the given IP address is valid. Values 0 to 59 seconds — Specifies the number of seconds that the given IP address is valid. Values 0 to 59

server

Syntax	server <i>server1</i> [<i>server2</i>]
Context	<pre>config>service>vprn>if>dhcp config>service>vprn>nw-if>dhcp config>service>vprn>sub-if>grp-if>dhcp</pre>
Description	This command specifies a list of servers where requests will be forwarded. The list of servers can entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list. The flood command is applicable only in the VPLS case. There is a scenario with VPLS where the VPLS node only wants to add Option 82 information to the DHCP request to provider per-subscriber information, but it does not do full DHCP relay. In this case, the server is set to “flood”. This means the DHCP request is still a broadcast and is sent through the VPLS domain. A node running at L3 further upstream then can perform the full L3 DHCP relay function.
Default	no server
Parameters	server — Specifies the DHCP server IP address. There can be a maximum of 8 DHCP servers configured.

python-policy

Syntax	python-policy <i>name</i> no python-policy

Context	config>service>vprn>if>dhcp
Description	This command specifies a python policy to be used for DHCPv4. Python policies are configured in the config>python> python-policy <i>name</i> context.
Parameters	<i>name</i> — Specifies the name of an existing python script up to 32 characters in length.

python-policy

Syntax	python-policy <i>name</i> no python-policy
Context	config>service>vprn>if>dhcp6-relay
Description	This command specifies a python policy to be used for DHCPv6 relay. Python policies are configured in the config>python> python-policy <i>name</i> context.
Parameters	<i>name</i> — Specifies the name of an existing python script up to 32 characters in length.

relay-proxy

Syntax	relay-proxy [release-update-src-ip] [siaddr-override <i>ip-address</i>] no relay-proxy
Context	config>service>vprn>if>dhcp config>service>vprn>sub-if>dhcp config>service>vprn>sub-if>grp-if>dhcp
Description	This command enables the DHCPv4 relay proxy function on the interface. The command has no effect when no dhcp servers are configured (DHCPv4 relay not configured). By default, unicast DHCPv4 release messages are forwarded transparently.

A relay proxy enhances the relay such that it also relays unicast client DHCPv4 REQUEST messages (lease renewals).

- In the upstream direction, update the source IP address and add the gateway IP address (gi-address) field before sending the message to the intended DHCP server (the message is not broadcast to all configured DHCP servers).
- In the downstream direction, remove the gi-address and update the destination IP address to the address of the *yiaddr* (your IP address) field.

The optional **release-update-src-ip** parameter updates the source IP address of a DHCP RELEASE message with the address used for relayed DHCPv4 messages.

The optional **siaddr-override** *ip-address* parameter enables DHCP server IP address hiding towards the client. This parameter requires that **lease-populate** is enabled on the interface. The DHCP server ip address is required for the address hiding function and is stored in the lease state record. The client interacts with the relay proxy as if it is the DHCP server. In all DHCP messages to the client, the value of following header fields and DHCP options containing the DHCP server IP address is replaced with the configured *<ip-address>*:

- the “source IP address” field in the IP DHCPv4 packet header
- the “siaddr” field in the DHCPv4 header if not equal to zero in the message received from the server
- the Server Identification option (DHCPv4 option 54) if present in the original server message
- the source IP address field in the IP packet header

DHCP OFFER selection during initial binding is done in the relay-proxy. Only the first DHCP OFFER message is forwarded to the client. Subsequent DHCP OFFER messages from different servers are silently dropped.

Default	no relay-proxy
Parameters	release-update-src-ip — Updates the source IP address of a DHCP RELEASE message with the address used for relayed DHCPv4 messages ip-address — Enables DHCPv4 server address hiding towards the DHCPv4 client and activates DHCPv4 OFFER selection in case multiple DHCP servers are configured. The <i>ip-address</i> can be any local address in the same routing instance. If DHCP relay lease-split is enabled, siaddr-override ip-address has priority over the emulated-server ip-address configured in the proxy-server and will be used as the source IP address.

relay-plain-bootp

Syntax	[no] relay-plain-bootp
Context	config>service>vprn>if>dhcp
Description	This command enables the relaying of plain BOOTP packets. The no form of the command disables the relaying of plain BOOTP packets.

snoop

Syntax	[no] snoop
Context	config>service>vprn>nw-if>dhcp
Description	This command enables snooping of DHCP packets on this interface.

The **no** form of the command disables snooping.

trusted

Syntax	[no] trusted
Context	config>service>vprn>if>dhcp config>service>vprn>nw-if>dhcp config>service>vprn>sub-if>grp-if>dhcp
Description	According to RFC 3046, <i>DHCP Relay Agent Information Option</i>, a DHCP request where the giaddr is 0.0.0.0 and which contains an Option 82 field in the packet, should be discarded, unless it arrives on a “trusted” circuit. If trusted mode is enabled on an IP interface, the relay agent router will modify the request's giaddr to be equal to the ingress interface and forward the request. This behavior only applies when the action in the Relay Agent Information Option is “keep”. In the case where the Option 82 field is being replaced by the relay agent (action = replace), the original Option 82 information is lost anyway, and there is thus no reason for enabling the trusted option. The no form of this command returns the system to the default.
Default	no trusted

egress

Syntax	egress
Context	config>service>vprn>nw-if
Description	This command enters the context to configure egress network filter policies for the interface.

use-arp

Syntax	[no] use-arp
Context	config>service>vprn>if>dhcp
Description	This command enables the use of ARP to determine the destination hardware address. The no form of the command disables the use of ARP to determine the destination hardware address

user-db

Syntax	user-db <i>local-user-db-name</i> no user-db
Context	config>service>vprn>sub-if>grp-if>dhcp
Description	This command configures the local user database to use for authentication. The no form of the command removes the value from the configuration.
Default	no user-db
Parameters	<i>local-user-db-name</i> — Specifies the local user database to use for authentication.

dynamic-tunnel-redundant-next-hop

Syntax	dynamic-tunnel-redundant-next-hop <i>ip-address</i> no dynamic-tunnel-redundant-next-hop
Context	config>service>vprn>if
Description	This command specifies redundant next-hop address on public or private IPsec interface (with public or private tunnel-sap) for dynamic IPsec tunnel. The specified next-hop address will be used by standby node to shunt traffic to master in case of it receives them. The next-hop address will be resolved in routing table of corresponding service.
Parameters	<i>ip-address</i> — Specifies the dynamic ISA tunnel redundant next-hop address.

egr-ip-load-balancing

Syntax	egr-ip-load-balancing { source destination inner-ip } no egr-ip-load-balancing
Context	config>service>vprn>if>load-balancing config>service>vprn>if>nw-if>load-balancing
Description	This command specifies whether to include source address or destination address or both in LAG/ECMP hash on IP interfaces. Additionally, when I4-load-balancing is enabled the command applies also to inclusion of source/destination port in the hash inputs. The no form of this command includes both source and destination parameters.
Default	no egr-ip-load-balancing
Parameters	<i>source</i> — Specifies using source address and (if I4-load balancing is enabled) source port in the hash, ignore destination address/port.

destination — Specifies using destination address and (if I4-load balancing is enabled) destination port in the hash, ignore source address/port.

inner-ip — Specifies use of the inner IP header parameters instead of outer IP header parameters in LAG/ECMP hash for IPv4 encapsulated traffic.

enable-ingress-stats

Syntax	[no] enable-ingress-stats
Context	config>router>interface config>service>ies >interface config>service>vprn>interface config>service>ies>sub-if>grp-if config>service>vprn>sub-if>grp-if
Description	This command enables the collection of ingress interface IP stats. This command is only applicable to IP statistics, and not to uRPF statistics. If enabled, then the following statistics are collected: • IPv4 offered packets • IPv4 offered octets • IPv6 offered packets • IPv6 offered octets • Octet statistics for IPv4 and IPv6 bytes at IP interfaces include the layer 2 frame overhead.
Default	no enable-ingress-stats

enable-mac-accounting

Syntax	[no] enable-mac-accounting
Context	config>service>vprn>if
Description	This command enables MAC accounting functionality on this interface. The no form of the command disables MAC accounting functionality on this interface.

host-connectivity-verify

Syntax	host-connectivity-verify [source {vrrp interface}] [interval <i>interval</i>] [action {remove alarm}] [timeout <i>retry-timeout</i>] [retry-count <i>count</i>] host-connectivity-verify [interval <i>interval</i>] [action {remove alarm}] [timeout <i>retry-timeout</i>] [retry-count <i>count</i>] [family <i>family</i>]

Context	<pre>config>service>vprn>if config>service>vprn>sub-if>grp-if config>service>vprn>sub-if>grp-if>dhcp</pre>
Description	This command enables subscriber host connectivity verification on a given SAP within a service. This tool will periodically scan all known hosts (from dhcp-state) and perform a UC ARP request. The subscriber host connectivity verification will maintain state (connected vs. not-connected) for all hosts.
Default	no host-connectivity-verify
Parameters	source {vrrp interface} — Specifies the source to be used for generation of subscriber host connectivity verification packets. The interface keyword forces the use of the interface mac and ip addresses. There are up to 16 possible subnets on a given interface, therefore subscriber host connectivity verification tool will use always an address of the subnet to which the given host is pertaining. In case of group-interfaces. one of the parent subscriber-interface subnets (depending on host's address) will be used. interval interval — Specifies the interval, expressed in minutes, which specifies the time interval which all known sources should be verified. The actual rate is then dependent on number of known hosts and interval. Values 1 to 6000 (A zero value can be used by the SNMP agent to disable host-connectivity-verify.) action {remove alarm} — Defines the action taken on a subscriber host connectivity verification failure for a given host. The remove keyword raises an alarm and removes dhcp-state and releases all allocated resources (queues, table entries, and so on). DHCP-RELEASE will be signaled to corresponding DHCP server. Static hosts will never be removed. The alarm keyword raises an alarm indicating that the host is disconnected. timeout retry-timeout — Specifies the timeout in seconds between consecutive retries of subscriber host connectivity verification checks, in case the host does not respond. Values 10 to 60 seconds retry-count count — Specifies the number of retries that will be carried out before a subscriber host is considered to have failed the SHCV check. Values 2 to 29 family family — Indicates the IP address family for which subscriber host connectivity verification checks will be enabled. It can be set to ipv4 or ipv6 only, or both. Values 2 to 29

hold-time

Syntax **hold-time**

Context config>service>vprn>interface
config>service>vprn>network-interface
config>service>vprn>subscriber-interface
config>service>vprn>redundant-interface

Description This command creates the CLI context to configure interface level hold-up and hold-down timers for the associated IP interface.

The **up** timer controls a delay for the associated IPv4 or IPv6 interface so that the system will delay the deactivation of the associated interface for the specified amount of time.

The **down** timer controls a delay for the associated IPv4 or IPv6 interface so that the system will delay the activation of the associated interface for the specified amount of time

up

Syntax up ip *seconds*
no up ip
up ipv6 *seconds*
no up ipv6

Context config>service>vprn>if>hold-time
config>service>vprn>nw-if>hold-time
config>service>vprn>sub-if>hold-time
config>service>vprn>red-if>hold-time

Description This command will cause a delay in the deactivation of the associated IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface down.

The **no** form of the command removes the command from the active configuration and removes the delay in deactivating the associated IP interface. If the configuration is removed during a delay period, the currently running delay will continue until it expires.

Parameters *seconds* — The time delay, in seconds, to make the interface operational.

Values 1 to 1200

down

Syntax down ip *seconds* [init-only]
no up ip
up ipv6 *seconds* [init-only]
no up ipv6

Context config>service>vprn>if>hold-time
config>service>vprn>nw-if>hold-time

	<pre>config>service>vprn>sub-if>hold-time config>service>vprn>red-if>hold-time</pre>
Description	This command will cause a delay in the activation of the associated IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface up, unless the init-only option is configured. If the init-only option is configured, the delay is only applied when the IP interface is first configured or after a system reboot. The no form of the command removes the command from the active configuration and removes the delay in activating the associated IP interface. If the configuration is removed during a delay period, the currently running delay will continue until it completes.
Parameters	<i>seconds</i> — The time delay, in seconds, to make the interface operational. Values 1 to 1200 init-only — Specifies that the down delay is only applied when the interface is configured or after a reboot. Values 1 to 1200

3.8.2.19 Interface ICMP Commands

icmp

Syntax	icmp
Context	<pre>config>service>vprn>if config>service>vprn>sub-if>grp-if config>service>vprn>nw-if</pre>
Description	This command configures Internet Control Message Protocol (ICMP) parameters on a VPRN service.

mask-reply

Syntax	[no] mask-reply
Context	<pre>config>service>vprn>if>icmp config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp</pre>
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 will reply to mask requests.

The **no** form of this command disables replies to ICMP mask requests on the router interface.

Default mask-reply — Specifies to reply to ICMP mask requests.

packet-too-big

Syntax	packet-too-big <i>number seconds</i> no packet-too-big
Context	config>service>vprn>if>ipv6>icmp6
Description	This command configures the rate for Internet Control Message Protocol version 6 (ICMPv6) packet-too-big messages.
Parameters	<i>number</i> — Specifies the number of packet-too-big messages to send in the time frame specified by the <i>seconds</i> parameter. Values 10 to 1000 Default 100 <i>seconds</i> — Specifies the time frame, in seconds, that is used to limit the number of packet-too-big messages issued. Values 1 to 60 Default 10

param-problem

Syntax	param-problem <i>number seconds</i> no param-problem [<i>number seconds</i>]
Context	config>service>vprn>if>icmp config>service>vprn>if>ipv6>icmp6 config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
Description	This command specifies whether parameter-problem ICMP messages should be sent. When enabled, parameter-problem ICMP messages are generated by this interface. The no form of the command disables the sending of parameter-problem ICMP messages.

Parameters	<i>number</i> — Specifies the number of parameter-problem ICMP messages to send in the time frame specified by the <i>seconds</i> parameter. Values 10 to 1000 Default 100 <i>seconds</i> — Specifies the time frame, in seconds, that is used to limit the number of parameter-problem ICMP messages issued. Values 1 to 60 Default 10

redirects

Syntax	redirects [<i>number seconds</i>] no redirects
Context	config>service>vprn>if>icmp config>service>vprn>if>ipv6>icmp6 config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
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 <i>number</i> and <i>seconds</i> parameters, by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval. By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 messages 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 — Specifies a maximum of 100 redirect messages in 10 seconds.
Parameters	<i>number</i> — Specifies the maximum number of ICMP redirect messages to send. This parameter must be specified with the <i>seconds</i> parameter. Values 10 to 1000 <i>seconds</i> — Specifies the time frame in seconds used to limit the <i>seconds</i> of ICMP redirect messages that can be issued. Values 1 to 60

time-exceeded

Syntax	time-exceeded <i>number seconds</i> no time-exceeded
Context	config>service>vprn>if>ipv6>icmp6
Description	This command configures rate for ICMPv6 time-exceeded messages.
Parameters	<i>number</i> — Specifies the maximum number of time-exceeded messages to send, expressed as a decimal integer. This parameter must be specified with the <i>seconds</i> parameter. Values 10 to 1000 <i>seconds</i> — Specifies the time frame in seconds used to limit the <i>number</i> of time-exceeded messages that can be issued, expressed as a decimal integer. Values 1 to 60

ttl-expired

Syntax	ttl-expired <i>number seconds</i> no ttl-expired
Context	config>service>vprn>if>icmp config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
Description	This command configures the rate of Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface. By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval. The no form of this command disables the limiting the rate of TTL expired messages on the router interface.
Default	ttl-expired 100 10
Parameters	<i>number</i> — Specifies the maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the <i>seconds</i> parameter. Values 10 to 2000 <i>seconds</i> — Specifies the time frame in seconds used to limit the <i>number</i> of ICMP TTL expired messages that can be issued, expressed as a decimal integer. Values 1 to 60

unreachables

Syntax	unreachables [<i>number seconds</i>] no unreachables
Context	config>service>vprn>if>icmp config>service>vprn>if>ipv6>icmp6 config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
Description	This command enables and configures the rate for ICMP host and network destination unreachable messages issued on the router interface. The unreachables command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional <i>number</i> and <i>seconds</i> parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval. By default, generation of ICMP destination unreachable messages is enabled at a maximum rate of 10 messages 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	<i>number</i> — Specifies the maximum number of ICMP unreachable messages to send. This parameter must be specified with the <i>seconds</i> parameter. Values 10 to 2000 <i>seconds</i> — Specifies the time frame in seconds used to limit the <i>number</i> of ICMP unreachable messages that can be issued. Values 1 to 60

ip-mtu

Syntax	ip-mtu <i>octets</i> no ip-mtu
Context	config>service>vprn>nw-if
Description	This command configures the IP maximum transmit unit (packet) for the associated router IP interface. The configured IP-MTU cannot be larger than the calculated IP MTU based on the port MTU configuration. The MTU that is advertised from the IES size is:

MINIMUM((SdpOperPathMtu - EtherHeaderSize), (Configured ip-mtu))

The **no** form of the command returns the associated IP interfaces MTU to its default value, which is calculated based on the port MTU setting. For Ethernet ports this will typically be 1554.

Default	no ip-mtu
Parameters	<i>octets</i> — Specifies the octets.
Values	512 to 9000

lag

Syntax	lag <i>lag-id[:encap-val]</i> no lag				
Context	config>service>vprn>nw-if				
Description	This command binds the interface to a Link Aggregation Group (LAG) The no form of the command removes the LAG id from the configuration.				
Parameters	<i>lag-id[:encap-val]</i> — Specifies the LAG ID.				
Values	<table> <tr> <td>lag-id</td><td>1 to 800</td></tr> <tr> <td>encap-val</td><td>0 (for null) 0 to 4094 (for dot1q)</td></tr> </table>	lag-id	1 to 800	encap-val	0 (for null) 0 to 4094 (for dot1q)
lag-id	1 to 800				
encap-val	0 (for null) 0 to 4094 (for dot1q)				

lag-per-link-hash

Syntax	lag-per-link-hash class {1 2 3} weight [1 to 1024] no per-link-hash
Context	config>service>vprn>nw-if
Description	This command configures weight and class to this SAP to be used on LAG egress when the LAG uses weighted per-link-hash. The no form of this command restores the default configuration.
Default	no lag-per-link-hash (equivalent to weight 1 class 1)

3.8.2.20 Interface SAP ATM Commands

atm

Syntax	atm
Context	config>service>vprn>if>sap config>service>vprn>sub-if>grp-if>sap
Description	This command enters the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supporting ATM functionality such as: • Configuring ATM port or ATM port-related functionality on MDAs supporting ATM functionality. • Configuring ATM-related configuration for ATM-based SAPs that exist on MDAs supporting ATM functionality. If ATM functionality is not supported for a given context, the command returns an error.

egress

Syntax	egress
Context	config>service>vprn>if>sap>atm config>service>vprn>sub-if>grp-if>sap>atm
Description	This command configures egress ATM attributes for the SAP.

encapsulation

Syntax	encapsulation <i>atm-encap-type</i>
Context	config>service>vprn>if>sap>atm config>service>vprn>sub-if>grp-if>sap>atm
Description	This command configures RFC 2684, <i>Multiprotocol Encapsulation over ATM AAL5</i>, encapsulation for an ATM PVCC delimited SAP. This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition also references the ATM Forum LAN Emulation specification. The encapsulation is driven by the services for which the SAP is configured. Ingress traffic that does not match the configured encapsulation will be dropped.
Default	encapsulation aal5snap-routed (for VPRN service SAPs)

Parameters	<i>atm-encap-type</i> — specifies the encapsulation type
Values	aal5snap-routed — Routed encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684. aal5mux-ip — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684. aal5snap-bridged — Bridged encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684. aal5mux-bridged-eth-nofcs — Bridged IP encapsulation for VC multiplexed circuit as defined in RFC 2684.

ingress

Syntax	ingress
Context	config>service>vprn>if>sap>atm config>service>vprn>sub-if>grp-if>sap>atm
Description	This command configures ingress ATM attributes for the SAP.

traffic-desc

Syntax	traffic-desc <i>traffic-desc-profile-id</i> no traffic-desc
Context	config>service>vprn>if>sap>atm>egress config>service>vprn>if>sap>atm>ingress config>service>vprn>sub-if>grp-if>sap>atm>egress config>service>vprn>sub-if>grp-if>sap>atm>ingress
Description	This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP). When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction. When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction. The no form of the command reverts the traffic descriptor to the default traffic descriptor profile.
Default	The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.
Parameters	<i>traffic-desc-profile-id</i> — Specifies a defined traffic descriptor profile (see the QoS atm-td-profile command).

oam

Syntax	oam
Context	config>service>vprn>if >sap>atm config>service>vprn>sub-if>grp-if>sap>atm
Description	This command enters the context to configure OAM functionality for a PVCC delimiting a SAP. The ATM-capable MDAs support F5 end-to-end OAM functionality (AIS, RDI, Loopback): • ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95 • GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996 • GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

alarm-cells

Syntax	[no] alarm-cells
Context	config>service>vprn>if>sap>atm>oam config>service>vprn>sub-if>grp-if>sap>atm>oam
Description	This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC termination to monitor and report the status of their connection by propagating fault information through the network and by driving PVCC's operational status. When alarm-cells functionality is enabled, a PVCC's operational status is affected when a PVCC goes into an AIS or RDI state because of an AIS/RDI processing (assuming nothing else affects PVCC's operational status, for example, if the PVCC goes DOWN, or enters a fault state and comes back UP, or exits that fault state). RDI cells are generated when PVCC is operationally DOWN. No OAM-specific SNMP trap is raised whenever an endpoint enters/exits an AIS or RDI state, however, if as result of an OAM state change, the PVCC changes operational status, then a trap is expected from an entity the PVCC is associated with (for example a SAP). The no command disables alarm-cells functionality for a PVCC. When alarm-cells functionality is disabled, a PVCC's operational status is no longer affected by a PVCC's OAM state changes due to AIS/RDI processing (when alarm-cells is disabled, a PVCC will change operational status to UP due to alarm-cell processing) and RDI cells are not generated as result of the PVCC going into AIS or RDI state. The PVCC's OAM status, however, will record OAM faults as described above.
Default	enabled for PVCCs delimiting VPRN SAPs

periodic-loopback

Syntax	[no] periodic-loopback
Context	config>service>vprn>if >sap>atm>oam config>service>vprn>sub-if>grp-if>sap>atm
Description	This command enables periodic OAM loopbacks on this SAP. This command is only configurable on VPRN SAPs. When enabled, an ATM OAM loopback cell is transmitted every period as configured in the config>system>atm>oam>loopback-period <i>period</i> context. If a response is not received and consecutive retry-down retries also result in failure, the endpoint will transition to an alarm indication signal/loss of clock state. Then, an ATM OAM loopback cell will be transmitted every period as configured in the loopback-period <i>period</i>. If a response is received for the periodic loopback and consecutive retry-up retries also each receive a response, the endpoint will transition back to the up state. The no form of the command sets the value back to the default.
Default	no periodic-loopback

3.8.2.21 Interface Anti-Spoofing Commands

anti-spoof

Syntax	anti-spoof {ip mac ip-mac nh-mac} no anti-spoof-type
Context	config>service>vprn>if>sap config>service>vprn>sub-if>grp-if>sap
Description	This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the interface. The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (ip, mac, ip-mac, nh-mac) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled. The no form of the command reverts to the default.
Default	Filter type default types: • anti-spoof ip (Non-Ethernet encapsulated SAP) • anti-spoof ip-mac (Ethernet encapsulated SAP) • no anti-spoof-type (other SAPs)

-
- Parameters**
- ip** — Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the anti-spoof type **ip** command fails.
 - mac** — Configures SAP anti-spoof filtering to use only the source MAC address in its lookup. Setting the anti-spoof filter type to **mac** is not allowed on non-Ethernet encapsulated SAPs. If a static host exists on the SAP without a specified MAC address, the anti-spoof type **mac** command fails. The anti-spoof type **mac** command will also fail if the SAP does not support Ethernet encapsulation.
 - ip-mac** — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup. If a static host exists on the SAP without both the IP address and MAC address specified, the anti-spoof type **ip-mac** command fails. This is also true if the default anti-spoof filter type of the SAP is **ip-mac** and the default is not overridden. The anti-spoof type **ip-mac** command will also fail if the SAP does not support Ethernet encapsulation.
 - nh-mac** — Indicates that the ingress anti-spoof is based on the source MAC address and the egress anti-spoof is based on the nh-ip-address.

app-profile

- Syntax** **app-profile** *app-profile-name*
no app-profile
- Context** config>service>vprn>if>sap
 config>service>vprn>sub-if>grp-if>sap
- Description** This command configures the application profile name.
- Parameters** *app-profile-name* — Specifies an existing application profile name configured in the **config>app-assure>group>policy** context.

arp-limit

- Syntax** **arp-limit** *limit* [**log-only**] [**threshold** *percent*]
no arp-limit
- Context** config>service>vprn>interface
- Description** This command configures the maximum amount of dynamic IPv4 ARP entries that can be learned on an IP interface.
- When the number of dynamic ARP entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations will be dropped. Entries that have already been learned will be refreshed.
- The **no** form of the command removes the **arp-limit**.

Default	90 percent
Parameters	log-only — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit will be learned. <i>percent</i> — The threshold value (as a percentage) that triggers a warning message to be sent. Values 0 to 100 <i>limit</i> — The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic ARP learning is disabled and no dynamic ARP entries are learned. Values 0 to 524288

arp-populate

Syntax	[no] arp-populate
Context	config>service>vprn>if config>service>vprn>sub-if>subscriber-interface config>service>vprn>sub-if>grp-if
Description	This command enables populating static and dynamic hosts into the system ARP cache. When enabled, the host's IP address and MAC address are placed in the system ARP cache as a managed entry. Static hosts must be defined on the interface using the host command. Dynamic hosts are enabled on the system through enabling lease-populate in the IP interface DHCP context. In the event that both a static host and a dynamic host share the same IP and MAC address, the system's ARP cache retains the host information until both the static and dynamic information are removed. Both static and dynamic hosts override static ARP entries. Static ARP entries are marked as inactive when they conflict with static or dynamic hosts and will be repopulated once all static and dynamic host information for the IP address are removed. Since static ARP entries are not possible when static subscriber hosts are defined or when DHCP lease state table population is enabled, conflict between static ARP entries and the arp-populate function is not an issue. The arp-populate command fails if an existing static subscriber host on the SAP does not have both MAC and IP addresses specified. Once arp-populate is enabled, creating a static subscriber host on the SAP without both an IP address and MAC address fails. arp-populate can only be enabled on VPRN interfaces supporting Ethernet encapsulation. Use the no form of the command to disable ARP cache population functions for static and dynamic hosts on the interface. All static and dynamic host information in the systems ARP cache will be removed. Any existing static ARP entries previously inactive due to static or dynamic hosts will be populated in the system ARP cache.

When **arp-populate** is enabled, the system will not send out ARP Requests for hosts that are not in the ARP cache. Only statically configured and DHCP learned hosts are reachable through an IP interface with arp-populate enabled.

Default no arp-populate

arp-retry-timer

Syntax **arp-retry-timer** *timer-multiple*
no arp-retry-timer

Context config>service>vprn>if
config>service>vprn>network-interface

Description This command allows the arp retry timer to be configured to a specific value.

The timer value is entered as a multiple of 100 ms. So a timer value of 1, means the ARP timer will be set to 100 ms.

The **no** form of this command removes the command from the active configuration and returns the ARP retry timer to its default value of 5 s.

Default arp-retry-timer 50

Parameters *timer-multiple* — Specifies the multiple of 100 ms that the ARP retry timer will be configured as.

Values 1 to 300 (equally a timer range of 100 ms to 30 000 ms)

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>vprn>if
config>service>vprn>sub-if>grp-if

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 restores **arp-timeout** to the default value.

Default arp-timeout 14400

Parameters	<i>seconds</i> — 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

authentication-policy

Syntax	authentication-policy <i>name</i> no authentication-policy
Context	config>service>vprn>if config>service>vprn>sub-if>grp-if
Description	This command assigns an authentication policy to the interface. The no form of this command removes the policy name from the group interface configuration.
Default	no authentication-policy
Parameters	<i>name</i> — Specifies the authentication policy name. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

bandwidth

Syntax	bandwidth <i>bandwidth</i> no bandwidth
Context	config>service>vprn>if>sap
Description	This command specifies the admin bandwidth assigned to SAPs, ports and LAGs which is used by SAP bandwidth CAC. SAP: Attempts to increase the SAP admin bandwidth will fail if there is insufficient available admin bandwidth on its port or LAG, otherwise the port or LAG available admin bandwidth will be reduced by the incremental SAP admin bandwidth. Reducing the SAP admin bandwidth will increase the available admin bandwidth on its port or LAG. This is not supported for PW-SAPs, Ethernet tunnels or subscriber group interface SAPs. The no version of the command reverts to the default value.
Default	no bandwidth
Parameters	<i>bandwidth</i> — Specifies the admin bandwidth assigned to the SAP, port or LAG, in kb/s. Values 1 to 3200000000

calling-station-id

Syntax	calling-station-id <i>calling-station-id</i> no calling-station-id
Context	config>service>vprn>sub-if>grp-if>sap config>service>vprn>if>sap
Description	This command enables the inclusion of the calling-station-id attribute in RADIUS authentication requests and RADIUS accounting messages. The value inserted is set at the SAP level. If no value is set at the SAP level, an empty string is included.
Default	no calling-station-id

host

Syntax	[no] host {[ip <i>ip-address</i> [mac <i>ieee-address</i>]} [subscriber <i>sub-ident-string</i>] [sub-profile <i>sub-profile-name</i>] [sla-profile <i>sla-profile-name</i>] no host {[ip <i>ip-address</i>] [mac <i>ieee-address</i>]}
Context	config>service>vprn>if>sap
Description	This command creates a static host for the SAP. Applications within the system that make use of static host entries include anti-spoof, and source MAC population into the VPLS forwarding database. Multiple static hosts can be defined on the SAP. Each host is identified by a source IP address, a source MAC address, or both a source IP and source MAC address. When anti-spoof is enabled on the SAP, the host information will be populated into the SAP's anti-spoof table, allowing ingress packets matching the entry access to the SAP. When the MAC address exists in the host definition, the MAC address is populated into the VPLS forwarding database and associates it with the SAP. The static host definition overrides any static MAC entries using the same MAC and prevents dynamic learning of the MAC on another interface. Defining a static host identical to an existing static host has no effect and will not generate a log or error message. Every static host definition must have at least one address defined, IP or MAC. Static hosts may exist on the SAP even with anti-spoof and arp-populate (VPRN) features disabled. When enabled, each feature has different requirements for static hosts.
Default	There are no default static entries.

-
- Parameters**
- anti-spoof** — When enabled, this feature uses static and dynamic host information to populate entries into an anti-spoof filter table. The anti-spoof filter entries generated will be of the same type as specified in the anti-spoof type parameter. If the SAP anti-spoof filter is defined as *mac*, each static host definition must specify a MAC address. If the SAP anti-spoof filter is defined as *ip*, each static host definition must specify an IP address. If the SAP anti-spoof filter is defined as *ip-mac*, each static host definition must specify both an IP address and MAC address. If definition of a static host is attempted without the appropriate addresses specified for the enabled anti-spoof filter, the static host definition fails.
- arp-populate** — When enabled, this feature uses static and dynamic host information to populate entries into the system's ARP cache. This is only available on the VPRN service SAPs. Both a MAC address and IP address are required to populate an ARP entry in the system. If definition of a static host is attempted without both a MAC and IP address specified when *arp -populate* is enabled, the static host definition fails.
- fdb-populate** — This is an implicit feature that uses the static host definition as a static MAC in the VPLS forwarding database. It cannot be enabled or disabled and has no effect on the ability to create static hosts without a MAC address specified. When a MAC address is specified for a static host, it will automatically be populated into the VPLS forwarding database associated with the SAP on which the host is created. The static host MAC address will override any static MAC entries using the same MAC and prevent dynamic learning of the MAC on another interface. Existing static MAC entries with the same MAC address as a static host are marked as inactive but not deleted. If all static hosts are removed from the SAP, the static MAC may be populated. New static MAC definitions for the VPLS instance may be created while a static host exists associated with the static MAC address.
- The **no** form of the command removes a static entry from the system. The specified **ip address** and **mac address** must match the host's exact IP and MAC addresses as defined when it was created. When a static host is removed from the SAP, the effect of its removal on the anti-spoof filter, ARP cache or the VPLS forwarding database is also evaluated.
- ip ip-address** — Specifies this optional parameter when defining a static host. The IP address must be specified for **anti-spoof ip** and **anti-spoof ip-mac** commands. Only one static host can be configured on the SAP with a given IP address.
- The following rules apply to configure static hosts using an IP address:
- Only one static host can be defined using a specific IP address.
 - Defining a static host with the same IP address as a previous static host overwrites the previous static host.
 - If a static host has an IP address assigned, the MAC address for the host is optional (depending on the features enabled on the SAP).
- mac mac-address** — Specifies this optional parameter when defining a static host. The MAC address must be specified for **anti-spoof mac**, and **anti-spoof ip-mac**. Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address. The following rules apply to configuring static hosts using a MAC address:
- Multiple static hosts may share the same MAC address.

- Executing the host command with the same MAC address but a different IP address as an existing static host will create a new static host.
- If a static host has a MAC address assigned, the IP address for the host is optional (depending on the features enabled on the SAP).

Values 8k static and dynamic hosts per 10G forwarding complex. 64k8k per system.

subscriber *sub-ident-string* — This optional parameter specifies an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

- For VPRN SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's *sub-ident-string* is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPRN destinations.

If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.

If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

ARP requests are never forwarded back to the same SAP or within the receiving SAP's split horizon group.

sub-profile *sub-profile-name* — Specifies this optional parameter to specify an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

sla-profile *sla-profile-name* — Specifies this optional parameter to specify an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

frame-relay

Syntax	frame-relay
Context	config>service>vprn>if>sap
Description	This command enters the context to configure Frame Relay parameters on the SAP.

frf-12

Syntax	[no] frf-12
Context	config>service>vprn>if>sap
Description	This command enables the use of FRF12 headers. The no form of the command disables the use of FRF12 headers.

ete-fragment-threshold

Syntax	ete-fragment-threshold <i>threshold</i> no ete-fragment-threshold				
Context	config>service>vprn>if>sap>frf-12				
Description	This command specifies the maximum length of a fragment to be transmitted. The no form of the command reverts to the default.				
Parameters	<i>threshold</i> — Specifies the maximum length of a fragment to be transmitted. <table> <tr> <td>Values</td><td>128 to 512</td></tr> <tr> <td>Default</td><td>0</td></tr> </table>	Values	128 to 512	Default	0
Values	128 to 512				
Default	0				

interleave

Syntax	[no] interleave
Context	config>service>vprn>if>sap>frame-relay>frf.12
Description	This command enables interleaving of high priority frames and low-priority frame fragments within a FR SAP using FRF.12 end-to-end fragmentation. When this option is enabled, only frames of the FR SAP non expedited forwarding class queues are subject to fragmentation. The frames of the FR SAP expedited queues are interleaved, with no fragmentation header, among the fragmented frames. In effect, this provides a behavior like in MLPPP Link Fragment Interleaving (LFI). When this option is disabled, frames of all the FR SAP forwarding class queues are subject to fragmentation. The fragmentation header is however not included when the frame size is smaller than the user configured fragmentation size. In this mode, the SAP transmits all fragments of a frame before sending the next full or fragmented frame. The receive direction of the FR SAP supports both modes of operation concurrently, with and without fragment interleaving. The no form of this command restores the default mode of operation.

Default no interleave

scheduling-class

Syntax **scheduling-class** *class-id*

Context config>service>vprn>if>sap

Description This command specifies the scheduling class to use for this SAP.

Parameters *class-id* — Specifies the scheduling class to use for this SAP.

Values 0 to 3

Default 0

host-lockout-policy

Syntax **host-lockout-policy** *policy-name*
 no host-lockout-policy

Context config>service>vprn>if>sap

Description This command configures a host lockout policy.

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

host-shutdown

Syntax **[no] host-shutdown**

Context config>service>vprn>if>sap

Description This command administratively enables host creation on this SAP.

ip-tunnel

Syntax **ip-tunnel** *name* [**create**]
 no ip-tunnel *name*

Context config>service>vprn>if>sap

Description This command is used to configure an IP-GRE or IP-IP tunnel and associate it with a private tunnel SAP within an IES or VPRN service.

The **no** form of the command deletes the specified IP/GRE or IP-IP tunnel from the configuration. The tunnel must be administratively shut down before issuing the **no ip-tunnel** command.

- Default** No IP tunnels are defined.
- Parameters** *ip-tunnel-name* — Specifies the name of the IP tunnel. Tunnel 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 between double quotes.

backup-remote-ip

- Syntax** **backup-remote-ip** *ip-address*
no backup-remote-ip
- Context** config>service>vprn>if>sap>ip-tunnel
- Description** This command sets the backup destination IPv4 address of GRE encapsulated packets associated with a particular GRE tunnel. If the primary destination address is not reachable in the delivery service (there is no route) or not defined then this is the destination IPv4 address of GRE encapsulated packets sent by the delivery service.
- The **no** form of the command deletes the backup-destination address from the GRE tunnel configuration.
- Parameters** *ip-address* — Specifies the destination IPv4 address of the GRE tunnel.
- Values** 1.0.0.0 to 223.255.255.255

delivery-service

- Syntax** **delivery-service** *service-id*
delivery-service name *service-name*
no delivery-service
- Context** config>service>vprn>if>sap>ip-tunnel
- Description** This command sets the delivery service for GRE encapsulated packets associated with a particular GRE tunnel. This is the IES or VPRN service where the GRE encapsulated packets are injected and terminated. The delivery service may be the same service that owns the private tunnel SAP associated with the GRE tunnel. The GRE tunnel does not come up until a valid delivery service is configured.
- The **no** form of the command deletes the delivery-service from the GRE tunnel configuration.
- Parameters** *service-id* — identifies the service used to originate and terminate the GRE encapsulated packets belonging to the GRE tunnel.

This variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **delivery-service name** *service-name* variant can be used in all configuration modes.

Values {*id* | *svc-name*}

id: 1 to 2147483647

svc-name: up to 64 characters (*svc-name* is an alias for input only. The *svc-name* gets replaced with an id automatically by SR OS in the configuration).

service-name — identifies the service used to originate and terminate the GRE encapsulated packets belonging to the GRE tunnel.

Values 1 to 64 characters

dscp

Syntax	dscp <i>dscp-name</i> no dscp
Context	config>service>vprn>if>sap>ip-tunnel
Description	This command sets the DSCP code-point in the outer IP header of GRE encapsulated packets associated with a particular GRE tunnel. The default, set using the no form of the command, is to copy the DSCP value from the inner IP header (after remarking by the private tunnel SAP egress qos policy) to the outer IP header.
Default	no dscp
Parameters	<i>dscp</i> — Specifies the DSCP code-point to be used.
Values	be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

private-tcp-mss-adjust

Syntax	private-tcp-mss-adjust <i>octets</i> no private-tcp-mss-adjust
Context	config>service>vprn>if>sap>ipsec-tunnel config>service>vprn>if>sap>ip-tunnel

Description	This command enables TCP MSS adjust for L2TPv3 tunnels on the private side of the service level. When this command is configured, the system updates the TCP MSS option value of the received TCP SYN packet on the private side. Note that this command can be overridden by the corresponding configuration on the group or tunnel level. The no form of this command disables TCP MSS adjust on the private side.
Default	no private-tcp-mss-adjust
Parameters	<i>octets</i> — Specifies the new TCP MSS value in octets. Values 512 to 9000

public-tcp-mss-adjust

Syntax	public-tcp-mss-adjust <i>octets</i> no public-tcp-mss-adjust
Context	config>service>vprn>if>sap>ipsec-tunnel config>service>vprn>if>sap>ip-tunnel
Description	This command enables TCP MSS adjust for L2TPv3 tunnels on the public side on the service level. When the command is configured, the system updates the TCP MSS option value of the received TCP SYN packet on the public side that is encapsulated in the L2TPv3 tunnel. Note that this command can be overridden by the corresponding configuration on the group or tunnel level. The no form of this command disables TCP MSS adjust on the public side.
Default	no public-tcp-mss-adjust
Parameters	<i>octets</i> — Specifies the new TCP MSS value in octets Values 512 to 9000

remote-ip

Syntax	remote-ip <i>ip-address</i> no remote-ip
Context	config>service>vprn>if>sap>ip-tunnel
Description	This command sets the primary destination IPv4 address of GRE encapsulated packets associated with a particular GRE tunnel. If this address is reachable in the delivery service (there is a route) then this is the destination IPv4 address of GRE encapsulated packets sent by the delivery service.

The **no** form of the command deletes the destination address from the GRE tunnel configuration.

Parameters *ip-address* — Specifies the destination IPv4 address of the GRE tunnel.
Values 1.0.0.0 to 223.255.255.255

source

Syntax **source** *ip-address*
no source

Context config>service>vprn>if>sap>ip-tunnel

Description This command sets the source IPv4 address of GRE encapsulated packets associated with a particular GRE tunnel. It must be an address in the subnet of the associated public tunnel SAP interface. The GRE tunnel does not come up until a valid source address is configured.

The **no** form of the command deletes the source address from the GRE tunnel configuration. The tunnel must be administratively shut down before issuing the **no source** command.

Parameters *ip-address* — Specifies the source IPv4 address of the GRE tunnel.
Values 1.0.0.0 to 223.255.255.255

3.8.2.22 Interface SAP Filter and QoS Policy Commands

egress

Syntax **egress**

Context config>service>vprn>if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command enters the context to configure egress SAP Quality of Service (QoS) policies and filter policies.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing. If no egress filter is defined, no filtering is performed.

ingress

Syntax **ingress**

Context config>service>vprn>if>sap

```
config>service>vprn>sub-if>grp-if>sap
```

Description This command enters the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

agg-rate

Syntax [no] **agg-rate**

Context config>service>vprn>if>sap>egress
config>service>vprn>sub-if>grp-if>sap>egress

Description This command is used to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, and **queue-frame-based-accounting**.

limit-unused-bandwidth

Syntax [no] **limit-unused-bandwidth**

Context config>service>vprn>if>sap>egress>agg-rate
config>service>vprn>sub-if>grp-if>sap>egress

Description This command is used to enable (or disable) aggregate rate overrun protection on the agg-rate context.

queue-frame-based-accounting

Syntax [no] **queue-frame-based-accounting**

Context config>service>vprn>if>sap>egress>agg-rate
config>service>vprn>sub-if>grp-if>sap>egress

Description This command is used to enabled (or disable) frame based accounting on all policers and queues associated with the agg-rate context. Only supported on Ethernet ports. Not supported on HSMDA Ethernet ports. Packet byte offset settings are not included in the applied rate when queue frame-based accounting is configured; the offsets are applied to the statistics.

rate

Syntax **rate {max | rate}**

no rate

Context config>service>vprn>if>sap>egress>agg-rate
config>service>vprn>sub-if>grp-if>sap>egress

Description This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the agg-rate context to be considered to be active on the context's object (SAP, subscriber, Vport, and so on).

agg-rate-limit

Syntax **agg-rate-limit** *agg-rate* [**queue-frame-based-accounting**]
no agg-rate-limit

Context config>service>vprn>sub-if>grp-if>wlan-gw>egress

Description This command defines a maximum total rate for all egress queues on a service SAP or multi-service site. The **agg-rate-limit** command is mutually exclusive with the egress scheduler policy. When an egress scheduler policy is defined, the **agg-rate-limit** command fails. If the **agg-rate-limit** command is specified, an attempt to bind a **scheduler-policy** to the SAP or multi-service site fails.

A multi-service site must have a port scope defined that ensures all queues associated with the site are on the same port or channel. If the scope is not set to a port, the agg-rate-limit command fails. Once an agg-rate-limit has been assigned to a multi-service site, the scope cannot be changed to card level.

A port scheduler policy must be applied on the egress port or channel the SAP or multi-service site is bound to in order for the defined agg-rate-limit to take effect. The egress port scheduler enforces the aggregate queue rate as it distributes its bandwidth at the various port priority levels. The port scheduler stops offering bandwidth to member queues once it has detected that the aggregate rate limit has been reached.

If a port scheduler is not defined on the egress port, the queues are allowed to operate based on their own bandwidth parameters.

The **no** form of the command removes the aggregate rate limit from the SAP or multi-service site.

Parameters *agg-rate* — Defines the rate, in kilobits per second, that the maximum aggregate rate that the queues on the SAP or multi-service site can operate.

Values 1 to 40000000, **max**

queue-frame-based-accounting — This keyword enables frame based accounting on all queues associated with the SAP or Multi-Service Site. If frame based accounting is required when an aggregate limit is not necessary, the max keyword should precede the queue-frame-based-accounting keyword. If frame based accounting must be disabled, execute agg-rate-limit without the queue-frame-based-accounting keyword present.

Default Frame based accounting is disabled by default.

filter

Syntax	filter ip <i>ip-filter-id</i> no filter
Context	config>service>vprn>if>sap>egress config>service>vprn>if>sap>ingress config>service>vprn>sub-if>grp-if>sap>egress
Description	This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface. Filter policies control the forwarding and dropping of packets based on IP matching criteria. The filter command is used to associate a filter policy with a specified <i>ip-filter-id</i> with an ingress or egress SAP. The <i>ip-filter-id</i> must already be defined before the filter command is executed. If the filter policy does not exist, the operation fails and an error message returned. In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets. The no form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use scope command within the filter definition to change the scope to local or global. The default scope of a filter is local.
Parameters	ip <i>ip-filter-id</i> — Specifies IP filter policy. The filter ID must already exist within the created IP filters.
Values	1 to 65535

policer-control-policy

Syntax	policer-control-policy <i>policy-name</i> no policer-control-policy
Context	config>service>vprn>sub-if>grp-if>sap>egress
Description	This command is used to specify a policer control policy to apply to SAP egress.

Parameters *policy-name* — Specifies the name of a policer control policy. 32 characters maximum.

hsmdda-queue-override

Syntax **[no] hsmdda-queue-override**

Context config>service>vprn>if>sap>egress

Description This command enters the context to configure HSMDDA queue overrides.

queue

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

Context config>service>vprn>if>sap>egress>hsmdda-queue-override

Description This command configures overrides for a HSMDDA queue. The actual valid values are those defined in the given SAP QoS policy.

Parameters *queue-id* — Specifies the queue ID to override.

Values 1 to 8

create — This keyword is mandatory when creating a new queue override.

packet-byte-offset

Syntax **packet-byte-offset {add add-bytes | subtract sub-bytes}**
 no packet-byte-offset

Context config>service>vprn>if>sap>egress>hsmdda-queue-override

Description This command adds or subtracts the specified number of bytes to the accounting function for each packet handled by the HSMDDA queue. Normally, the accounting and leaky bucket functions are based on the Ethernet DLC header, payload and the 4 byte CRC (everything except the preamble and inter-frame gap). As an example, the packet-byte-offset command can be used to add the frame encapsulation overhead (20 bytes) to the queues accounting functions.

The accounting functions affected include:

- Offered High Priority / In-Profile Octet Counter
- Offered Low Priority / Out-of-Profile Octet Counter
- Discarded High Priority / In-Profile Octet Counter
- Discarded Low Priority / Out-of-Profile Octet Counter

- Forwarded In-Profile Octet Counter
- Forwarded Out-of-Profile Octet Counter
- Peak Information Rate (PIR) Leaky Bucket Updates
- Committed Information Rate (CIR) Leaky Bucket Updates
- Queue Group Aggregate Rate Limit Leaky Bucket Updates

The secondary shaper leaky bucket, scheduler priority level leaky bucket and the port maximum rate updates are not affected by the configured packet-byte-offset. Each of these accounting functions are frame based and always include the preamble, DLC header, payload and the CRC regardless of the configured byte offset.

The packet-byte-offset command accepts either add or subtract as valid keywords which define whether bytes are being added or removed from each packet traversing the queue. Up to 31 bytes may be added to the packet and up to 32 bytes may be removed from the packet. An example use case for subtracting bytes from each packet is an IP based accounting function. Given a Dot1Q encapsulation, the command packet-byte-offset subtract 14 would remove the DLC header and the Dot1Q header from the size of each packet for accounting functions only. The 14 bytes are not actually removed from the packet, only the accounting size of the packet is affected.

As inferred above, the variable accounting size offered by the packet-byte-offset command is targeted at the queue and queue group level. The packet-byte-offset, when set, applies to all queues in the queue group. The accounting size of the packet is ignored by the secondary shapers, the scheduling priority level shapers and the scheduler maximum rate. The actual on-the-wire frame size is used for these functions to allow an accurate representation of the behavior of the subscriber's packets on an Ethernet aggregation network.

The packet-byte-offset value may be overridden at the queue-group level.

Parameters **add** *add-bytes* — Indicates that the byte value should be added to the packet for queue and queue group level accounting functions. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command. The **add** keyword is mutually exclusive with the **subtract** keyword.

Values 0 to 31

subtract *sub-bytes* — Indicates that the byte value should be subtracted from the packet for queue and queue group level accounting functions. The **subtract** keyword is mutually exclusive with the **add** keyword. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command.

Values 1 to 64

slope-policy

Syntax **slope-policy** *hsmda-slope-policy-name*
no slope-policy

Context	config>service>vprn>if>sap>egress>hsmda-queue-override
Description	This command specifies an existing slope policy name.

wrr-weight

Syntax	wrr-weight value no wrr-weight
Context	config>service>vprn>if>sap>egress>hsmda-queue-override>queue
Description	This command assigns the weight value to the HSMDA queue. The no form of the command returns the weight value for the queue to the default value.
Parameters	<i>percentage</i> — Specifies the weight for the HSMDA queue. Values 1to 32

wrr-policy

Syntax	wrr-policy hsmda-wrr-policy-name no wrr-policy
Context	config>service>vprn>if>sap>egress>hsmda-queue-override
Description	This command associates an existing HSMDA weighted-round-robin (WRR) scheduling loop policy to the HSMDA queue.
Parameters	<i>hsmda-wrr-policy-name</i> — Specifies the existing HSMDA WRR policy name to associate to the queue.

secondary-shaper

Syntax	secondary-shaper secondary-shaper-name no secondary-shaper
Context	config>service>vprn>if>sap>egress>hsmda-queue-override
Description	This command configures an HSMDA secondary shaper. A shaper override can only be configured on an HSMDA SAP.
Parameters	<i>secondary-shaper-name</i> — Specifies a secondary shaper name up to 32 characters in length.

policer-override

Syntax	[no] policer-override
Context	config>service>vprn>if>sap>egress config>service>vprn>if>sap>ingress
Description	This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to one or more policers created on the SAP through the sap-ingress or sap-egress QoS policies. The no form of the command is used to remove any existing policer overrides.
Default	no policer-override

policer

Syntax	policer <i>policer-id</i> [create] no policer <i>policer-id</i>
Context	config>service>vprn>if>sap>egress>policer-override config>service>vprn>if>sap>ingress>policer-override
Description	This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to a specific policer created on the SAP through a sap-ingress or sap-egress QoS policy. The no form of the command is used to remove any existing overrides for the specified policer-id.
Parameters	<i>policer-id</i> — This parameter is required when executing the policer command within the policer-override context. The specified <i>policer-id</i> must exist within the sap-ingress or sap-egress QoS policy applied to the SAP. If the policer is not currently used by any forwarding class or forwarding type mappings, the policer will not actually exist on the SAP. This does not preclude creating an override context for the <i>policer-id</i>. create — The create keyword is required when a policer 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 configuration, the create keyword is not required.

match-qinq-dot1p

Syntax	match-qinq-dot1p {top bottom} no match-qinq-dot1p
Context	config>service>vprn>if>sap>ingress config>service>vprn>sub-if>grp-if>sap>ingress

Description This command specifies which Dot1Q tag position Dot1P bits in a QinQ encapsulated packet should be used to evaluate Dot1P QoS classification.

The **match-qinq-dot1p** command allows the top or bottom PBits to be used when evaluating the applied sap-ingress QoS policy's Dot1P entries. The **top** and **bottom** keywords specify which position should be evaluated for QinQ encapsulated packets.

The **no** form of the command restores the default dot1p evaluation behavior for the SAP.

By default, the bottom most service delineating Dot1Q tags Dot1P bits are used. [Table 41](#) defines the default behavior for Dot1P evaluation when the **match-qinq-dot1p** command is not executed.

Table 41 Dot1P Default Behavior

Port / SAP Type	Existing Packet Tags	PBits Used for Match
null	none	none
null	Dot1P (VLAN-ID 0)	Dot1P PBits
null	—	Dot1Q PBits
null	TopQ BottomQ	TopQ PBits
null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	none (Default SAP)	none
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

Default no match-qinq-dot1p - No filtering based on p-bits.

top or bottom must be specified to override the default QinQ dot1p behavior.

Parameters **top** — The top parameter is mutually exclusive to the bottom parameter. When the top parameter is specified, the top most PBits are used (if existing) to match any dot1p dot1p-value entries. [Table 42](#) defines the dot1p evaluation behavior when the top parameter is specified.

Table 42 Dot1P Evaluation Behavior

Port / SAP Type	Existing Packet Tags	PBits Used for Match
null	none	none
null	Dot1P (VLAN-ID 0)	Dot1P PBits
null	Dot1Q	Dot1Q PBits
null	TopQ BottomQ	TopQ PBits
null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	none (Default SAP)	none
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits

bottom — The bottom parameter is mutually exclusive to the top parameter. When the bottom parameter is specified, the bottom most PBits are used (if existing) to match any dot1p dot1p-value entries. The following tables define the bottom position QinQ and TopQ SAP dot1p evaluation and the default dot1p explicit marking actions.

Table 43 Bottom Position QinQ and TopQ SAP Dot1P Evaluation

Port / SAP Type	Existing Packet Tags	PBits Used for Match
null	none	none
null	Dot1P (VLAN-ID 0)	Dot1P PBits
null	Dot1Q	Dot1Q PBits
null	TopQ BottomQ	BottomQ PBits
null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	none (default SAP)	none
Dot1Q	Dot1P (default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	BottomQ PBits

Table 43 Bottom Position QinQ and TopQ SAP Dot1P Evaluation

Port / SAP Type	Existing Packet Tags	PBits Used for Match
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

Table 44 Default Dot1P Explicit Marking Actions

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
null	no preserved Dot1P bits	none
null	preserved Dot1P bits	preserved tag PBits remarked using dot1p-value
Dot1Q	no preserved Dot1P bits	new PBits marked using dot1p-value
Dot1Q	preserved Dot1P bits	preserved tag PBits remarked using dot1p-value
TopQ	no preserved Dot1P bits	TopQ PBits marked using dot1p-value
TopQ	preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits marked using dot1p-value, BottomQ PBits preserved
QinQ	no preserved Dot1P bits	TopQ PBits and BottomQ PBits marked using dot1p-value
QinQ	preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits and BottomQ PBits marked using dot1p-value

The dot1p dot1p-value command must be configured without the qinq-mark-top-only parameter to remove the TopQ PBits only marking restriction.

qinq-mark-top-only

Syntax [no] qinq-mark-top-only

Context config>service>vprn>if>sap>egress
config>service>vprn>sub-if>grp-if>sap>egress

Description	When enabled (the encapsulation type of the access port where this SAP is defined as qinq), the qinq-mark-top-only command specifies which P-bits/DEI bit to mark during packet egress. When disabled, both set of P-bits/DEI bit are marked. When the enabled, only the P-bits/DEI bit in the top Q-tag are marked.
Default	no qinq-mark-top-only

qos

Syntax	qos <i>policy-id</i> [port-redirect-group <i>queue-group-name</i> instance <i>instance-id</i>] no qos [<i>policy-id</i>]
Context	config>service>vprn>if>sap>egress config>service>vprn>sub-if>grp-if>sap>egress
Description	This command associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP). QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the policy- id does not exist, an error will be returned. The qos command is used to associate both ingress and egress QoS policies. The qos command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error. Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type will return an error. By default, no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used. The no form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.
Parameters	<i>policy-id</i> — The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist. Values 1 to 65535 <i>port-redirect-group</i> — This keyword associates a SAP egress with an instance of a named queue group template on the egress port of a given IOM/IMM/XMA. The queue-group-name and instance instance-id are mandatory parameters when executing the command. <i>queue-group-name</i> — Specifies the name of the egress port queue group of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid egress queue group, created under config>port>ethernet>access>egress.

instance *instance-id* — Specifies the instance of the named egress port queue group on the IOM/IMM/XMA.

Values 1 to 40960

Default 1

qos

Syntax	qos <i>policy-id</i> [shared-queuing multipoint-shared] qos <i>policy-id</i> [shared-queuing multipoint-shared] fp-redirect-group <i>queue-group-name</i> instance <i>instance-id</i> no qos [<i>policy-id</i>]
Context	config>service>vprn>if>sap>ingress config>service>vprn>sub-if>grp-if>sap>ingress
Description	This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP). QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the policy- id does not exist, an error will be returned. The qos command is used to associate both ingress and egress QoS policies. The qos command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error. Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type will return an error. By default, no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used. The no form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.
Parameters	<i>policy-id</i> — The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist. Values 1 to 65535 shared-queuing — Specifies the ingress shared queue policy used by this SAP. When the value of this object is null it means that the SAP will use individual ingress QoS queues instead of the shared ones. multipoint-shared — Specifies that this <i>queue-id</i> is for multipoint forwarded traffic only. This <i>queue-id</i> can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. Attempting to map forwarding class unicast traffic to a multipoint queue generates an error; no changes are made to the current unicast traffic queue mapping.

A queue must be created as multipoint. The **multipoint** designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the **multipoint** keyword, an error is generated and the command will not execute.

The **multipoint** keyword can be entered in the command line on a pre-existing multipoint queue to edit *queue-id* parameters.

Default Present (the queue is created as non-multipoint).

Values Multipoint or not present.

fp-redirect-group — Creates an instance of a named queue group template on the ingress forwarding plane of a given IOM/IMM/XMA. The queue-group-name and instance instance-id are mandatory parameters when executing the command. The named queue group template can contain only policers. If it contains queues, then the command fails.

queue-group-name — Specifies the name of the queue group template to be instantiated on the forwarding plane of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid ingress queue group template name, configured under *config>qos>queue-group-templates*.

instance-id — Specifies the instance of the named queue group to be created on the IOM/IMM/XMA ingress forwarding plane.

scheduler-policy

Syntax	scheduler-policy <i>scheduler-policy-name</i> no scheduler-policy
Context	config>service>vprn>if>sap>ingress config>service>vprn>if>sap>egress config>service>vprn>sub-if>grp-if>sap>egress config>service>vprn>sub-if>grp-if>sap>ingress
Description	This command applies an existing scheduler policy to an ingress or egress scheduler used by SAP policers and queues associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created once the customer site has been appropriately assigned to a chassis port, channel or slot. Scheduler policies are defined in the config>qos>scheduler-policy <i>scheduler-policy-name</i> context.

The **no** form of this command removes the configured ingress or egress scheduler policy from the multi-service customer site. When the policy is removed, the schedulers created due to the policy are removed also making them unavailable for the SAP policers and queues associated with the customer site. Policers and queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler. The SAPs that have policers and queues that rely on the removed schedulers enter into an operational state depicting the orphaned status of one or more policers and queues. When the **no scheduler-policy** command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.

Parameters	<i>scheduler-policy-name</i> : — The <i>scheduler-policy-name</i> parameter applies an existing scheduler policy that was created in the config>qos>scheduler-policy scheduler-policy-name context to create the hierarchy of ingress or egress virtual schedulers. The scheduler names defined within the policy are created and made available to any ingress or egress queues created on associated SAPs.
Values	any existing valid scheduler policy name

ipsec-gw

Syntax	ipsec-gw <i>name</i> no ipsec-gw
Context	config>service>vprn>if>sap
Description	This command configures the IPSec gateway.
Parameters	<i>name</i> — Specifies the IPSec gateway name up to 32 characters in length.

fwd-wholesale

Syntax	[no] fwd-wholesale
Context	config>service>vprn>sub-if>grp-if>sap
Description	This command enables the context to select specific protocols ingressing on the SAP to be redirected to another service. The command is applicable to static SAPs as well as PW-SAPs. The no form of this command removes the redirection.

pppoe

Syntax	pppoe <i>service-id</i> no pppoe
Context	config>service>vprn>sub-if>grp-if>sap>fwd-wholesale
Description	This command specifies that PPPoE packets on ingress on Ethertypes 0x8863 and 0x8864 will be redirected to the specified service. The service referred to by <i>svc-id</i> must be an Epipe service. Redirection to VC-switching Epipe services is not supported. The no form of the command removes the redirect.
Parameters	<i>service-id</i> — Specifies the service ID of the Epipe to which packets are redirected.
Values	1 to 2147483647 <i>svc-name</i> up to 64 characters

cert

Syntax	cert
Context	config>service>vprn>if>sap>ipsec-gw
Description	This command enters the context to configure certificate parameters.

cert-profile

Syntax	cert-profile <i>profile-name</i> no cert-profile
Context	config>service>vprn>if>sap>ipsec-gw>cert config>service>vprn>if>sap>ipsec-tun>dyn>cert
Description	This command specifies the cert-profile for the ipsec-tunnel or ipsec-gw. This command will override “cert” and “key” configuration under the ipsec-tunnel or ipsec-gw.
Parameters	<i>profile-name</i> — Specifies the name of cert-profile.

trust-anchor-profile

Syntax	trust-anchor <i>trust-anchor-profile-name</i> no trust-anchor-profile
Context	config>service>vprn>if>sap>ipsec-gw>cert
Description	This command configures the trust anchor profile name associated with this SAP IPSec tunnel certificate. This command will override “trust-anchor” configuration under the ipsec-tunnel or ipsec-gw.
Parameters	<i>trust-anchor-profile-name</i> — Specifies the name of trust-anchor-profile.

default-secure-service

Syntax	default-secure-service <i>service-id</i> interface <i>ip-int-name</i> default-secure-service name <i>service-name</i> interface <i>ip-int-name</i> no default-secure-service
Context	config>service>vprn>if>sap>ipsec-gw
Description	This command specifies a service ID or service name of the default security service used by this SAP IPSec gateway.
Parameters	<i>service-id</i> — Specifies a default secure service.

This variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **default-secure-service name** *service-name* variant can be used in all configuration modes.

Values {*id* | *svc-name*}

id: 1 to 2147483647

svc-name: up to 64 characters (*svc-name* is an alias for input only. The *svc-name* gets replaced with an id automatically by SR OS in the configuration).

service-name — Identifies the service, up to 64 characters.

default-tunnel-template

Syntax **default-tunnel-template** *ipsec template identifier*
no default-tunnel-template

Context config>service>vprn>if>sap>ipsec-gw

Description This command configures the default tunnel policy template for the gateway.

Parameters *ipsec template id** — [1 to 2048]

ike-policy

Syntax **ike-policy** *ike-policy-id*
no ike-policy

Context config>service>vprn>if>sap>ipsec-gw
config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying

Description This command configures the IKE policy for the gateway.

Parameters *ike-policy-id* — Specifies the IKE policy ID.

Values 1 to 2048

local-gateway-address

Syntax **local-gateway-address** *ip-address*
no local-gateway-address

Context config>service>vprn>if>sap>ipsec-gw

Description	This command configures the ipsec-gateway local address.
Parameters	<i>ip-address</i> — Specifies the IP unicast address.

local-id

Syntax	local-id type { ipv4 fqdn } [value [<i>value</i>]] no local-id
Context	config>service>vprn>if>sap>ipsec-gw
Description	This command specifies the local ID of the router used for IDi or IDr for IKEv2 tunnels. The local-id can only be changed or removed when tunnel or gateway is shut down. Default: Depends on local-auth-method such as: • Psk: local tunnel ip address • Cert-auth: subject of the local certificate
Parameters	type — Specifies the type of local ID payload, it could be ipv4 address/FQDN domain name/distinguish name of subject in X.509 certificate. Values ipv4 — Use ipv4 as the local ID type, the default value is the local tunnel end-point address. fqdn — Use FQDN as the local ID type, the value must be configured. dn — Use the subject of the certificate configured for the tunnel or gateway.

pre-shared-key

Syntax	pre-shared-key <i>key</i> no pre-shared-key
Context	config>service>vprn>if>sap>ipsec-gw config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying
Description	This command specifies the shared secret between the two peers forming the tunnel.
Parameters	<i>key</i> — Specifies a pre-shared-key for dynamic-keying.

radius-accounting-policy

Syntax	radius-accounting-policy <i>policy-name</i> no radius-accounting-policy

Context	config>service>vprn>if>sap>ipsec-gw config>service>vprn>l2tp config>service>vprn>l2tp>group config>service>vprn>l2tp>group>tunnel
Description	This command configures the radius-accounting-policy.
Parameters	<i>policy-name</i> — 32 characters maximum.

radius-authentication-policy

Syntax	radius-authentication-policy <i>name</i> no radius-authentication-policy
Context	config>service>vprn>if>sap>ipsec-gw
Description	This command configures the radius-authentication-policy.
Parameters	<i>name</i> — Specifies the policy name up to 32 characters in length.

lag-link-map-profile

Syntax	lag-link-map-profile <i>link-map-profile-id</i> no lag-link-map-profile
Context	config>service>vprn>if>sap config>service>vprn> sub-if>grp-if >sap
Description	This command assigns a pre-configured lag link map profile to a SAP/network interface configured on a LAG or a PW port that exists on a LAG. Once assigned/de-assigned, the SAP/network interface egress traffic will be re-hashed over LAG as required by the new configuration. The no form of this command reverts the SAP/network interface to use per-flow, service or link hash as configured for the service/LAG.
Default	no lag-link-map-profile
Parameters	<i>link-map-profile-id</i> — An integer from 1 to 64 that defines a unique lag link map profile on which the LAG the SAP/network interface exist.

multi-service-site

Syntax	multi-service-site <i>customer-site-name</i> no multi-service-site <i>customer-site-name</i>
Context	config>service>vprn>if>sap


```
config>service>vprn>sub-if>grp-if>sap
```

Description

This command creates a new customer site or edits an existing customer site with the *customer-site-name* parameter. A customer site is an anchor point to create an ingress and egress virtual scheduler hierarchy. When a site is created, it must be assigned to a chassis slot or port on the 7750 SR. When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).

The scheduler policy association with the customer site normally prevents the scheduler policy from being deleted until after the scheduler policy is removed from the customer site. The multi-service-site object will generate a log message indicating that the association was deleted due to scheduler policy removal.

When the multi-service customer site is created, an ingress and egress scheduler policy association does not exist. This does not prevent the site from being assigned to a chassis slot or prevent service SAP assignment. After the site has been created, the ingress and egress scheduler policy associations can be assigned or removed at any time.

Parameters

customer-site-name — Each customer site must have a unique name within the context of the customer. If *customer-site-name* already exists for the customer ID, the CLI context changes to that site name for the purpose of editing the site scheduler policies or assignment. Any modifications made to an existing site will affect all SAPs associated with the site. Changing a scheduler policy association may cause new schedulers to be created and existing policers and queues on the SAPs to no longer be orphaned. Existing schedulers on the site may cease to exist, causing policers and queues relying on that scheduler to be orphaned.

If the *customer-site-name* does not exist, it is assumed that an attempt is being made to create a site of that name in the customer ID context. The success of the command execution depends on the following:

- The maximum number of customer sites defined for the chassis slot has not been met.
- The *customer-site-name* is valid.
- The **create** keyword is included in the command line syntax (if the system requires it).

When the maximum number of customer sites has been exceeded a configuration error occurs; the command will not execute and the CLI context will not change.

If the *customer-site-name* is invalid, a syntax error occurs; the command will not execute and the CLI context will not change.

Values

Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

static-host

Syntax	static-host ip <i>ip/did-address</i> [mac <i>ieee-address</i>] [create] static-host mac <i>ieee-address</i> [create] no static-host [ip <i>ip-address</i> >] mac <i>ieee-address</i> > no static-host all [force] no static-host ip <i>ip-address</i>
Context	config>service>vprn>if>sap config>service>vprn>sub-if>grp-if>sap
Description	This command configures a static host on this SAP.
Parameters	ip <i>ip-address</i> — Specifies the IPv4 unicast address. <i>mac ieee-address</i> — Specifies this optional parameter when defining a static host. Every static host definition must have at least one address defined, IP or MAC. force — Specifies the forced removal of the static host addresses. <i>sla-profile sla-profile-name</i> This optional parameter is used to specify an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the config>subscr-mgmt>sla-profile context.

ancp-string

Syntax	ancp-string <i>ancp-string</i> no ancp-string
Context	config>service>vprn>if>sap>static-host config>service>vprn>sub-if>grp-if>sap>static-host
Description	This command specifies the ANCP string associated to this SAP host.
Parameters	<i>ancp-string</i> — Specifies the ANCP string up to 63 characters in length.

app-profile

Syntax	app-profile <i>app-profile-name</i> [scope <i>scope-type</i>] no app-profile
Context	config>service>vprn>sub-if>grp-if>sap>static-host
Description	This command specifies an application profile name.
Parameters	<i>app-profile-name</i> — Specifies the application profile name, up to 32 characters in length.

scope-type — Specifies the scope to which the application profile is assigned, in the **config>service>vprn>sub-if>grp-if>sap>static-host** context.

Values subscriber—The application profile applies to this context with subscriber scope (all hosts or devices).
 mac—The application profile applies to this context with MAC scope (single device).

Default subscriber

inter-dest-id

Syntax **inter-dest-id** *intermediate-destination-id*
 no inter-dest-id

Context config>service>vprn>if>sap>static-host
 config>service>vprn>sub-if>grp-if>sap>static-host

Description This command specifies to which intermediate destination (for example a DSLAM) this host belongs.

Parameters *intermediate-destination-id* — Specifies the intermediate destination ID.

managed-routes

Syntax **managed-routes**

Context config>service>vprn>sub-if>grp-if>sap>static-host>managed-routes

Description This command configures managed routes.

route

Syntax **route** {*ip-prefix/length* | *ip-prefix netmask*} [**create**]
 no route {*ip-prefix/length* | *ip-prefix netmask*}

Context config>service>vprn>sub-if>grp-if>sap>static-host>managed-routes

Description This command assigns managed-route to a given subscriber-host. As a consequence, a static route pointing subscriber-host ip address as a next hop will be installed in FIB. Up to 16 managed routes per subscriber-host can be configured.

The **no** form of the command removes the respective route. Per default, there are no managed-routes configured.

sla-profile

Syntax	sla-profile <i>sla-profile-name</i> no sla-profile
Context	config>service>vprn>if>sap>static-host config>service>vprn>sub-if>grp-if>sap>static-host
Description	This command specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the config>subscr-mgmt>sla-profile context.
Parameters	<i>sla-profile-name</i> — Specifies the SLA profile name.

sub-profile

Syntax	sub-profile <i>sub-profile-name</i> no sub-profile
Context	config>service>vprn>if>sap>static-host config>service>vprn>sub-if>grp-if>sap>static-host
Description	This command specifies an existing subscriber profile name to be associated with the static subscriber host.
Parameters	<i>sub-profile-name</i> — Specifies the sub-profile name.

subscriber

Syntax	subscriber <i>sub-ident</i> no subscriber
Context	config>service>vprn>if>sap>static-host config>service>vprn>sub-if>grp-if>sap>static-host
Description	This command specifies an existing subscriber identification profile to be associated with the static subscriber host.
Parameters	<i>sub-ident</i> — Specifies the subscriber identification.

subscriber-sap-id

Syntax	[no] subscriber-sap-id
Context	config>service>vprn>if>sap>static-host config>service>vprn>sub-if>grp-if>sap>static-host

Description	This command enables using the SAP ID as subscriber id.
Parameters	subscriber-sap-id — Specifies to use the sap-id as the subscriber-id.

queue-group-redirect-list

Syntax	queue-group-redirect-list <i>redirect-list-name</i> no queue-group-redirect-list
Context	config>service>vprn>if>sap>egress config>service>vprn>if>sap>ingress
Description	This command applies a queue group redirect list to the ingress or egress of an interface SAP within a VPRN service. The redirect list is used to redirect traffic to different instances of the default queue group. This command requires the prior configuration of a default queue group instance, which is the queue group instance specified with the QoS policy under the SAP ingress or egress. The no version of this command removes the queue group redirect list from the SAP.
Parameters	<i>redirect-list-name</i> — Specifies the name of the queue group redirect list, up to 32 characters in length.

queue-override

Syntax	[no] queue-override
Context	config>service>vprn>if>sap>egress config>service>vprn>if>sap>ingress
Description	This command enters the context to configure override values for the specified SAP egress or ingress QoS queue. These values override the corresponding ones specified in the associated SAP egress or ingress QoS policy.

hs-secondary-shaper

Syntax	hs-secondary-shaper <i>policy-name</i> no hs-secondary-shaper
Context	config>service>vprn>if>sap>egress>queue-override
Description	This command configures the HS secondary shaper to be used to apply an aggregate rate and per-scheduling class rates to the SAP egress HSQ queue group.

The **no** form of the command removes the HS secondary shaper override from the configuration, returning the SAP egress HSQ queue group to the default HS secondary shaper on that port.

Parameters *policy-name* — Specifies the secondary shaper name, up to 32 characters.

hs-wrr-group

Syntax **hs-wrr-group** *group-id* [**create**]
no hs-wrr-group *group-id*

Context config>service>vprn>if>sap>egress>queue-override

Description This command configures the egress HS WRR group override parameters.
The **no** form of the command removes the group ID from the configuration.

Parameters *group-id* — Specifies the HS WRR group ID to override.

Values 1, 2

class-weight

Syntax **class-weight** *weight*
no class-weight

Context config>service>vprn>if>sap>egress>queue-override>hs-wrr-group

Description This command overrides the class weight of this WRR group at its parent primary shaper, relative to the other queues and WRR groups in different HSQ queue groups in the same scheduling class.
The **no** form of this command removes the class weight override value from the configuration.

Parameters *weight* — Specifies the class weight of the HS WRR group.

Values 1, 2, 4, 8

percent-rate

Syntax **percent-rate** *percent*
no percent-rate

Context config>service>vprn>if>sap>egress>queue-override>hs-wrr-group

Description	This command overrides the scheduling rate applied to the HS WRR group as a percentage of the port rate, including both the port's egress rate and port's HS scheduler policy max-rate, if configured. The override rate type must match the corresponding rate type within the applied QoS policy. The no form of this command removes the percent rate override value from the configuration.
Parameters	<i>percent</i> — Specifies the percent rate of the HS WRR group. Values 0.01 to 100.00

rate

Syntax	rate <i>rate</i> no <i>rate</i>
Context	config>service>vprn>if>sap>egress>queue-override>hs-wrr-group
Description	This command overrides the scheduling rate applied to the HS WRR group in Kb/s. Alternatively, the keyword max can be specified which removes the bandwidth limitation on the HS WRR group. The override rate type must match the corresponding rate type within the applied QoS policy. The no form of this command removes the rate override value from the configuration.
Parameters	<i>rate</i> — Specifies the scheduling rate of the HS WRR group in Kb/s. Values 1 to 2000000000, max

queue

Syntax	[no] <i>queue queue-id</i>
Context	config>service>vprn>if>sap>egress>queue-override config>service>vprn>if>sap>ingress>queue-override
Description	This command specifies the ID of the queue whose parameters are to be overridden.
Parameters	<i>queue-id</i> — The queue ID whose parameters are to be overridden. Values 1 to 32

adaptation-rule

Syntax	adaptation-rule [<i>pir adaptation-rule</i>] [<i>cir adaptation-rule</i>] no <i>adaptation-rule</i>
Context	config>service>vprn>if>sap>egress>queue-override>queue

```
config>service>vprn>if>sap>ingress>queue-override>queue
```

Description This command can be used to override specific attributes of the specified queue's adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the queue 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.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the adaptation rule is performed under the **hs-wrr-group** within the SAP egress QoS policy.

The **no** form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Default no adaptation-rule

Parameters

- pir* — The **pir** parameter defines the constraints enforced when adapting the PIR rate defined within the **queue queue-id rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.
- cir* — The **cir** parameter defines the constraints enforced when adapting the CIR rate defined within the **queue queue-id rate** 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 queue, while maintaining a minimum offset.

Values

- max** — The **max** (maximum) keyword is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.
- min** — The **min** (minimum) keyword is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.
- closest** — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

avg-frame-overhead

Syntax **avg-frame-overhead** *percent*
no avg-frame-overhead

Context	config>service>vprn>if>sap>egress>queue-override>queue
Description	This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a Sonet or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- **Offered-load** — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- **Frame encapsulation overhead** — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be 10000×0.1 or 1000 octets.

For egress Ethernet queues, the frame encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame encapsulation overhead would be 50×20 or 1000 octets.

- **Frame based offered-load** — The frame based offered-load is calculated by adding the offered-load to the frame encapsulation overhead. If the offered-load is 10000 octets and the encapsulation overhead was 1000 octets, the frame based offered-load would equal 11000 octets.
- **Packet to frame factor** — The packet to frame factor is calculated by dividing the frame encapsulation overhead by the queue's offered-load (packet based). If the frame encapsulation overhead is 1000 octets and the offered-load is 10000 octets then the packet to frame factor would be $1000 / 10000$ or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.
- **Frame based CIR** — The frame based CIR is calculated by multiplying the packet to frame factor with the queue's configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame based CIR would be 500×1.1 or 550 octets.
- **Frame based within-cir offered-load** — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.

- **Frame based PIR** — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be 7500 x 1.1 or 8250 octets.
- **Frame based within-pir offered-load** — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

Port scheduler operation using frame transformed rates — The port scheduler uses the frame based rates to determine the maximum rates that each queue may receive during the within-cir and above-cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

SAP and subscriber SLA-profile average frame overhead override — The average frame overhead parameter on a sap-egress may be overridden at an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers an avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

Default 0

Parameters *percent* — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

Values 0 to 100

burst-limit

Syntax **burst-limit** {default | size [bytes | kilobytes]}

no burst-limit

Context	config>service>vprn>sub-if>grp-if>sap>egress>queue-override>queue
Description	The queue burst-limit command defines an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate. The no form of this command restores the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies. When specified within a queue-override queue context, any current burst limit override for the queue is removed and the queue's burst limit is controlled by its defining policy.
Default	no burst-limit
Parameters	default — Reverts the queue's burst limit to the system default value. size — When a numeric value is specified (size), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and, by default, is interpreted as the burst limit in kilobytes. If the value is intended to be interpreted in bytes, the bytes qualifier must be added following size. Values 0 to 13671 kilobytes 0 to 14000000 bytes Default No default for size; use the default keyword to specify default burst limit. bytes — Specifies that the value given for size must be interpreted as the burst limit in bytes. kilobytes — Specifies that the value given for size must be interpreted as the burst limit in kilobytes. If neither bytes nor kilobytes is specified, the default qualifier is kilobytes.

cbs

Syntax	cbs size-in-kbytes no cbs
Context	config>service>vprn>if>sap>egress>queue-override>queue config>service>vprn>if>sap>ingress>queue-override>queue
Description	This command can be used to override specific attributes of the specified queue's CBS parameters. It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queue's CBS setting into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.

If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.

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

Default	no cbs
Parameters	<i>size-in-kbytes</i> — The size parameter is an integer expression of the number of kilobytes reserved for the queue. For a value of 10 kbytes, enter the number 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimum reserved size can be applied for scheduling purposes).
Values	0 to 131072 or default

drop-tail

Syntax	drop-tail
Context	config>service>vprn>if>sap>egress>queue-override>queue config>service>vprn>if>sap>ingress>queue-override>queue
Description	This command enters the context to configure queue drop tail parameters.

low

Syntax	low
Context	config>service>vprn>if>sap>egress>queue-override>queue>drop-tail config>service>vprn>if>sap>ingress>queue-override>queue>drop-tail
Description	This command enters the context to configure the queue low drop tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets are not accepted into the queue and will be discarded.

percent-reduction-from-mbs

Syntax	percent-reduction-from-mbs <i>percent</i> no percent-reduction-from-mbs
Context	config>service>vprn>if>sap>egress>queue-override>queue>drop-tail>low

config>service>vprn>if>sap>ingress>queue-override>queue>drop-tail>low

- Description** This command overrides the low queue drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and the percentage reduction is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes and out-of-profile packets will not be accepted into the queue if its depth is greater than this value, and so will be discarded.
- Parameters** *percent* — Specifies the percentage reduction from the MBS for a queue drop tail.
- Values** 0 to 100, default

hs-class-weight

- Syntax** **hs-class-weight** *weight*
no hs-class-weight
- Context** config>service>vprn>if>sap>egress>queue-override>queue
- Description** This command overrides the class weight of this queue at its parent primary shaper, relative to the other queues and WRR groups in different HSQ queue groups in the same scheduling class.
- The **no** form of this command removes the class weight override value from the configuration.
- Parameters** *weight* — Specifies the weight of the queue.
- Values** 1, 2, 4, 8

hs-wred-queue

- Syntax** **hs-wred-queue** *policy slope-policy-name*
no hs-wred-queue
- Context** config>service>vprn>if>sap>egress>queue-override>queue
- Description** This command overrides the slope policy applied to the HSQ queue group queue.
- The **no** form of this command removes the WRED queue policy override value from the configuration.
- Parameters** *slope-policy-name* — Specifies an existing slope policy name to apply to this HSQ queue group queue.

hs-wrr-weight

- Syntax** **hs-wrr-weight** *weight*

no hs-wrr-weight

Context	config>service>vprn>if>sap>egress>queue-override>queue
Description	This command overrides the WRR relative weight with which this queue should parent into an HSQ Weighted Round Robin (WRR) group defined within the associated HS attachment policy. The no form of this command removes the WRR weight override value from the configuration.
Parameters	<i>weight</i> — Specifies the HS WRR group queue weight.
Values	1 to 127

mbs

Syntax	mbs {size [bytes kilobytes] default} no mbs
Context	config>service>vprn>if>sap>egress>hsmda-queue-override>queue
Description	This command can be used to override specific attributes of the specified queue's MBS parameters. The MBS is a mechanism to override the default maximum size for the queue. The sum of the MBS for all queues on an egress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel. If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change. The no form of this command returns the MBS size assigned to the queue.
Default	mbs default
Parameters	<i>size</i> — Specifies the maximum number of kbytes of buffering allowed for the queue. For a value of 100 kb/s, enter the number 100. A value of 0 causes the queue to discard all packets. Values 0 to 2625 kilobytes 0 to 2688000 bytes default

mbs

Syntax	mbs { <i>size</i> [bytes kilobytes] default } no mbs
Context	config>service>vprn>if>sap>egress>queue-override>queue config>service>vprn>if>sap>ingress>queue-override>queue
Description	This command can be used to override specific attributes of the specified queue's MBS parameters. The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueueing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue. The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel. If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change. The no form of this command returns the MBS size assigned to the queue to the value.
Default	mbs default
Parameters	size — The size parameter is required when specifying mbs and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional bytes and kilobytes keywords are mutually exclusive and are used to explicitly define whether the size represents bytes or kilobytes. Values 0 to 1073741824 default bytes — When byte is defined, the value given for size is interpreted as the queue's MBS value given in bytes. kilobytes — When kilobytes is defined, the value is interpreted as the queue's MBS value given in kilobytes. default — Specifying the keyword default sets the MBS to its default value.

parent

Syntax	parent [weight <i>weight</i>] [cir-weight <i>cir-weight</i>] no parent
Context	config>service>vprn>if>sap>egress>queue-override>queue

Description	This command can be used to override the scheduler's parent weight and cir-weight information. The weights apply to the associated level/cir-level configured in the applied scheduler policy. The scheduler name must exist in the scheduler policy applied to the ingress or egress of the SAP or multi-service site. The override weights are ignored if the scheduler does not have a parent command configured in the scheduler policy – this allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the SAP/MSS overrides. If the parent scheduler does not exist causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non default weightings for fostered schedulers. The no form of the command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.
Default	no parent
Parameters	weight weight — Weight defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict level defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the queue or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit. A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level. Values 0 to 100 Default 1 cir-weight cir-weight — The cir-weight keyword defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same <i>cir-level</i> defined by the cir-level parameter in the applied scheduler policy. Within the strict cir-level, all cir-weight values from active children at that level are summed and the ratio of each active child's cir-weight to the total is used to distribute the available bandwidth at that level. A cir-weight is considered to be active when the policer, queue, or scheduler that the cir-weight pertains to has not reached the CIR and still has packets to transmit. A 0 (zero) cir-weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level. Values 0 to 100 Default 1

percent-rate

Syntax **percent-rate** *pir-percent* [**cir** *cir-percent*]

no percent-rate

Context	config>service>vprn>if>sap>egress>queue-override>queue
Description	The percent-rate command supports a queue's shaping rate and CIR rate as a percentage of the egress port's line rate. When the rates are expressed as a percentage within the template, the actual rate used per instance of the queue group queue-id will vary based on the port speed. For example, when the same template is used to create a queue group on a 1-Gigabit and a 10-Gigabit Ethernet port, the queue's rates will be 10 times greater on the 10 Gigabit port due to the difference in port speeds. This enables the same template to be used on multiple ports without needing to use port based queue overrides to modify a queue's rate to get the same relative performance from the queue. If the port's speed changes after the queue is created, the queue's shaping and CIR rates will be recalculated based on the defined percentage value. The rate and percent-rate commands override one another. If the current rate for a queue is defined using the percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at anytime. An egress port queue group queue rate override may be expressed as either a percentage or an explicit rate independent on how the queue's template rate is expressed. This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the percent-rate is performed under the hs-wrr-group within the SAP egress QoS policy. The no form of this command returns the queue to its default shaping rate and cir rate. When no percent-rate is defined within a port egress queue group queue override, the queue reverts to the defined shaping and CIR rates within the egress queue group template associated with the queue.
Parameters	<i>pir-percent</i> — Specifies the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate. Values 0.01 to 100.00 Default 100.00 <i>cir-percent</i> — Specifies the queue's committed scheduling rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate. Values 0.00 to 100.00 Default 100.00

rate

Syntax	rate <i>pir-rate</i> [cir <i>cir-rate</i>] no rate
Context	config>service>vprn>if>sap>egress>queue-override>queue config>service>vprn>if>sap>ingress>queue-override>queue
Description	This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters. The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth. The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop. The CIR can be used by the queue's parent commands <i>cir-level</i> and <i>cir-weight</i> parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler. The rate command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the <i>queue-id</i>. This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the rate is performed under the hs-wrr-group within the SAP egress QoS policy. The no form of the command returns all queues created with the <i>queue-id</i> by association with the QoS policy to the default PIR and CIR parameters (max, 0).
Default	rate max cir 0
Parameters	<i>pir-rate</i> — Defines the administrative PIR rate, in kilobits, for the queue. When the rate command is executed, a valid PIR setting must be explicitly defined. When the rate command has not been executed, the default PIR of max is assumed. Fractional values are not allowed and must be given as a positive integer. The actual PIR rate is dependent on the queue's adaptation-rule parameters and the actual hardware where the queue is provisioned. Values 1 to 2000000000, max Default max <i>cir-rate</i> — The cir parameter overrides the default administrative CIR used by the queue. When the rate command is executed, a CIR setting is optional. When the rate command has not been executed or the cir parameter is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and must be given as a positive integer. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues.

Values 0 to 2000000000, **max**

Default 0

rate

Syntax	rate <i>pir-rate</i> no rate
Context	config>service>vprn>if>sap>egress>hsmda-queue-override>queue
Description	This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR). The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth. The rate command can be executed at any time, altering the PIR rates for all queues created through the association of the SAP egress QoS policy with the <i>queue-id</i>. The no form of the command returns all queues created with the <i>queue-id</i> by association with the QoS policy to the default PIR parameters (max, 0).
Parameters	<i>pir-rate</i> — Defines the administrative PIR rate, in kilobits, for the queue. When the rate command is executed, a valid PIR setting must be explicitly defined. When the rate command has not been executed, the default PIR of max is assumed. Fractional values are not allowed and must be given as a positive integer. The actual PIR rate is dependent on the queue's adaptation-rule parameters and the actual hardware where the queue is provisioned. Values 1 to 100000000 Default max

scheduler-override

Syntax	[no] scheduler-override
Context	config>service>vprn>if>sap>egress config>service>vprn>if>sap>ingress
Description	This command specifies the set of attributes whose values have been overridden via management on this virtual scheduler. Clearing a given flag will return the corresponding overridden attribute to the value defined on the SAP's ingress scheduler policy.

scheduler

Syntax	scheduler <i>scheduler-name</i> no scheduler <i>scheduler-name</i>
Context	config>service>vprn>if>sap>egress>sched-override config>service>vprn>if>sap>ingress>sched-override
Description	This command can be used to override specific attributes of the specified scheduler name. A scheduler defines a bandwidth controls that limit each child (other schedulers, policers, and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have policers, queues, or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers. Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If <i>scheduler-name</i> already exists within the policy tier level (regardless of the inclusion of the keyword create), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs). If the <i>scheduler-name</i> exists within the policy on a different tier (regardless of the inclusion of the keyword create), an error occurs and the current CLI context will not change. If the <i>scheduler-name</i> does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following: Step 1. The maximum number of schedulers has not been configured. Step 2. The provided <i>scheduler-name</i> is valid. Step 3. The create keyword is entered with the command if the system is configured to require it (enabled in the environment create command). When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command will not execute, nor will the CLI context change. If the provided scheduler-name is invalid according to the criteria below, a name syntax error will occur, the command will not execute, and the CLI context will not change.
Parameters	<i>scheduler-name</i> — The name of the scheduler.
Values	Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

create — This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given *scheduler-name*. If the **create** keyword is omitted, **scheduler-name** is not created when the system environment variable *create* is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

parent

Syntax	parent [<i>weight weight</i>] [<i>cir-weight cir-weight</i>] no parent
Context	config>service>vprn>if>sap>ingress>sched-override>scheduler config>service>vprn>if>sap>egress>sched-override>scheduler
Description	This command can be used to override the scheduler's parent weight and cir-weight information. The weights apply to the associated level/cir-level configured in the applied scheduler policy. The scheduler name must exist in the scheduler policy applied to the ingress or egress of the SAP or multi-service site. The override weights are ignored if the scheduler does not have a parent command configured in the scheduler policy – this allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the SAP/MSS overrides. If the parent scheduler does not exist causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non default weightings for fostered schedulers. The no form of the command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.
Default	no parent
Parameters	weight weight — Weight defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict level defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the policer, queue, or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit. A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level. Values 0 to 100

cir-weight *cir-weight* — The **cir-weight** keyword defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same *cir-level* defined by the **cir-level** parameter in the applied scheduler policy. Within the strict **cir-level**, all **cir-weight** values from active children at that level are summed and the ratio of each active child's **cir-weight** to the total is used to distribute the available bandwidth at that level. A **cir-weight** is considered to be active when the policer, queue, or scheduler that the **cir-weight** pertains to has not reached the CIR and still has packets to transmit.

A 0 (zero) **cir-weight** value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict *cir-level*.

Values 0 to 100

rate

Syntax	rate <i>pir-rate</i> [cir <i>cir-rate</i>] no rate
Context	config>service>vprn>if>sap>egress>sched-override>scheduler config>service>vprn>if>sap>ingress>sched-override
Description	This command can be used to override specific attributes of the specified scheduler rate. The rate command defines the maximum bandwidth that the scheduler can offer its child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase. The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child policers and queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates. When a scheduler is defined without specifying a rate, the default rate is max. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child queues and schedulers to operate at their maximum rates. The no form of this command returns the scheduler's PIR and CIR parameters to the value configured in the applied scheduler policy.

Parameters	<i>pir-rate</i> — The pir parameter accepts a value of 1 to 3200000000, or the max keyword. Any other value will result in an error without modifying the current PIR rate. Values 1 to 3200000000, max <i>cir-rate</i> — The cir parameter accepts a value of 0 to 3200000000, or the max or sum keywords. Any other value will result in an error without modifying the current CIR rate. If the cir is set to max, then the CIR rate is set to infinity, but is limited by the <i>pir-rate</i>. If the cir is set to sum, then the CIR rate is set to the summed CIR values of the children schedulers, policers, or queues. Values 0 to 3200000000, max, sum

3.8.2.23 Interface VRRP Commands

vrrp

Syntax	vrrp <i>virtual-router-id</i> [owner] [passive] no vrrp <i>virtual-router-id</i>
Context	config>service>vprn>if
Description	This command creates or edits a Virtual Router ID (VRID) on the service IP interface. A VRID is internally represented in conjunction with the IP interface name. This allows the VRID to be used on multiple IP interfaces while representing different virtual router instances. Two VRRP nodes can be defined on an IP interface. One, both, or none may be defined as owner. The nodal context of vrrp <i>virtual-router-id</i> is used to define the configuration parameters for the VRID. The no form of this command removes the specified VRID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the VRID. The VRID does not need to be shut down in order to remove the virtual router instance.
Special Cases	Virtual Router Instance Owner IP Address Conditions — The virtual router instance owner can be created prior to assigning the parent IP interface primary or secondary IP addresses. In this case, the virtual router instance is not associated with an IP address. The operational state of the virtual router instance is down. VRRP Owner Command Exclusions — By specifying the VRRP <i>vrid</i> as owner, the following commands are no longer available: • vrrp priority — The virtual router instance owner is hard-coded with a priority value of 255 and cannot be changed. • vrrp master-int-inherit — Owner virtual router instances do not accept VRRP advertisement messages; the advertisement interval field is not evaluated and cannot be inherited.

- **ping-reply, telnet-reply and ssh-reply** — The **owner** virtual router instance always allows Ping, Telnet and SSH if the management and security parameters are configured to accept them on the parent IP interface.
- **vrrp shutdown** — The **owner** virtual router instance cannot be shut down on the **vrrp** node. If this was allowed, VRRP messages would not be sent, but the parent IP interface address would continue to respond to ARPs and forward IP packets. Another virtual router instance may detect the missing master due to the termination of VRRP advertisement messages and become master. This would result in two routers responding to ARP requests for the same IP addresses. To shut down the **owner** virtual router instance, use the **shutdown** command in the parent IP interface context. This will prevent VRRP participation, IP ARP reply and IP forwarding. To continue parent IP interface ARP reply and forwarding without VRRP participation, remove the **vrrp vrid** instance.
- **traceroute-reply**

VRRP Passive Command Exclusions — By specifying the VRRP *vrid* as **passive**, the following commands related to the master election and processing of VRRP advertisement messages are no longer available:

- **vrrp priority**
- **policy**
- **preempt**
- **master-int-inherit**
- **standby-forwarding**
- **int-delay**
- **message-interval**
- **authentication-key**
- **bfd-enable**

Parameters *virtual-router-id* — The virtual-router-id parameter specifies a new virtual router ID or one that can be modified on the IP interface.

Values 1 to 255

owner — Identifies this virtual router instance as owning the virtual router IP addresses. If the **owner** keyword is not specified at the time of *vrid* creation, the **vrrp backup** commands must be specified to define the virtual router IP addresses. The **owner** keyword is not required when entering the *vrid* for editing purposes. Once created as **owner**, a *vrid* on an IP interface cannot have the **owner** parameter removed. The *vrid* must be deleted, and then recreated without the **owner** keyword, to remove ownership.

passive — Identifies this virtual router instance as **passive**, and therefore, owning the virtual router IP addresses. A **passive vrid** does not send or receive VRRP advertisement messages, and is always in either the **master** state (if the interface is operational-up), or the **init** state (if the interface is operational-down). The **passive** keyword is not required when entering the *vrid* for editing purposes. Once a *vrid* on an IP interface is created as **passive**, the parameter cannot be removed from the *vrid*. The *vrid* must be deleted, and then recreated without the **passive** keyword, to remove parameter.

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>service>vprn>if>vrrp
Description	The authentication-key command, within the vrrp <i>virtual-router-id</i> context, is used to assign a simple text password authentication key to generate master VRRP advertisement messages and validate received VRRP advertisement messages.

The **authentication-key** command is one of the few commands not affected by the presence of the **owner** keyword. If simple text password authentication is not required, this command is not required. If the command is re-executed with a different password key defined, the new key will be used immediately. If a **no authentication-key** command is executed, the password authentication key is restored to the default value. The **authentication-key** command may be executed at any time.

To change the current in-use password key on multiple virtual router instances:

- Identify the current master
- Shut down the virtual router instance on all backups
- Execute the **authentication-key** command on the master to change the password key
- Execute the **authentication-key** command and the **no shutdown** command on each backup key

The **no** form of this command restores the default null string to the value of key.

Parameters	authentication-key — The <i>key</i> parameter identifies the simple text password used when VRRP Authentication Type 1 is enabled on the virtual router instance. Type 1 uses a string eight octets long that is inserted into all transmitted VRRP advertisement messages and compared against all received VRRP advertisement messages. The authentication data fields are used to transmit the key. The <i>key</i> parameter is expressed as a string consisting of up to eight alpha-numeric characters. Spaces must be contained in quotation marks (" "). The quotation marks are not considered part of the string.

The string is case sensitive and is left-justified in the VRRP advertisement message authentication data fields. The first field contains the first four characters with the first octet (starting with IETF RFC bit position 0) containing the first character. The second field holds the fifth through eighth characters. Any unspecified portion of the authentication data field is padded with the value 0 in the corresponding octet.

Values Any 7-bit printable ASCII character.

Exceptions:	Double quote (")	ASCII 34
	Carriage Return	ASCII 13
	Line Feed	ASCII 10
	Tab	ASCII 9
	Backspace	ASCII 8

hash-key — The hash key. The key can be any combination of ASCII characters up to 22 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ")

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

backup

Syntax	[no] backup <i>ip-address</i>
Context	config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp
Description	This command configures virtual router IP addresses for the interface.

bfd-enable

Syntax	[no] bfd-enable interface <i>interface-name</i> dst-ip <i>ip-address</i> [no] bfd-enable interface <i>interface-name</i> dst-ip <i>ip-address</i> name <i>name</i> [no] bfd-enable svc-id interface <i>interface-name</i> dst-ip <i>ip-address</i>

Context	<pre>config>service>vprn>if>vrrp config>service>vprn>sub-if>grp-if>srrp config>service>vprn>if>ipv6>vrrp</pre>
Description	This commands assigns a bi-directional forwarding (BFD) session providing heart-beat mechanism for the given VRRP/SRRP instance. There can be only one BFD session assigned to any given VRRP/SRRP instance, but there can be multiple SRRP/VRRP sessions using the same BFD session. If the interface used is configured with centralized BFD, the BFD transmit and receive intervals need to be set to at least 300 ms. BFD control the state of the associated interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface. The specified interface may not be configured with BFD; when it is, the virtual router will then initiate the BFD session. The no form of this command removes BFD from the configuration.
Parameters	svc-id — Specifies the service ID of the interface running BFD. If no svc-id is specified then it indicates that the interface is a network interface in the Base router instance. This variant of the command is only supported in 'classic' configuration-mode (configure system management-interface configuration-mode classic). The bfd-enable interface interface-name dst-ip ip-address name name variant can be used in all configuration modes. Values {<i>id</i> <i>svc-name</i>} <i>id</i>: 1 to 2147483647 <i>svc-name</i>: Specifies an existing service name up to 64 characters (<i>svc-name</i> is an alias for input only. The <i>svc-name</i> gets replaced with an id automatically by SR OS in the configuration) interface interface-name — Specifies the name of the interface running BFD, up to 32 characters. dst-ip ip-address — Specifies the destination address to be used for the BFD session. name name — Specifies a service name, up to 64 characters.

init-delay

Syntax	<pre>init-delay seconds no init-delay</pre>
Context	<pre>config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp</pre>
Description	This command configures a VRRP initialization delay timer.

Default	no init-delay
Parameters	<i>seconds</i> — Specifies the initialization delay timer for VRRP, in seconds.
Values	1 to 65535

mac

Syntax	[no] mac <i>ieee-mac-address</i>
Context	config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp
Description	This command assigns a specific MAC address to an IP interface. The no form of this command returns the MAC address of the IP interface to the default value.
Default	The physical MAC address associated with the Ethernet interface that the SAP is configured on.
Parameters	<i>ieee-mac-address</i> — Specifies the 48-bit MAC address for the static ARP in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> where <i>aa</i> , <i>bb</i> , <i>cc</i> , <i>dd</i> , <i>ee</i> and <i>ff</i> are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

master-int-inherit

Syntax	[no] master-int-inherit
Context	config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp
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 {[<i>seconds</i>] [milliseconds <i>milliseconds</i>]} no message-interval
Context	config>service>vprn>if config>service>vprn>if>ipv6>vrrp

Description	This command sets the advertisement timer and indirectly sets the master down timer on the virtual router instance. The message-interval setting must be the same for all virtual routers participating as a virtual router. Any VRRP advertisement message received with an Advertisement Interval field different than the virtual router instance configured message-interval value will be silently discarded. The message-interval command is available in both non-owner and owner vrrp <i>virtual-router-id</i> nodal contexts. If the message-interval command is not executed, the default message interval of 1 second will be used. The no form of this command restores the default message interval value of 1 second to the virtual router instance.
Parameters	seconds — The number of seconds that will transpire before the advertisement timer expires. Values 1 to 255 Default 1 milliseconds <i>milliseconds</i> — Specifies the milliseconds time interval between sending advertisement messages. This parameter is not supported on single-slot chassis. Values 100 to 900

oper-group

Syntax	oper-group <i>group-name</i> no oper-group
Context	config>service>vprn>if>vrrp
Description	This command configures VRRP to associate with an operational group. When associated, VRRP notifies the operational group of its state changes so that other protocols can monitor it to provide a redundancy mechanism. When VRRP is the master router (MR), the operational group is up and is down for all other VRRP states. The no form of the command removes the association.
Default	no oper-group — No operational group is configured.
Parameters	<i>group-name</i> — Specifies the operational group identifier up to 32 characters in length.

ping-reply

Syntax	[no] ping-reply
Context	config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp

Description	This command enables the non-owner master to reply to ICMP Echo Requests directed at the virtual router instances IP addresses. The ping request can be received on any routed interface. Ping must not have been disabled at the management security level (either on the parental IP interface or based on the Ping source host address). When ping-reply is not enabled, ICMP Echo Requests to non-owner master virtual IP addresses are silently discarded. Non-owner backup virtual routers never respond to ICMP Echo Requests regardless of the setting of ping-reply configuration. The ping-reply command is only available in non-owner vrrp <i>virtual-router-id</i> nodal context. If the ping-reply command is not executed, ICMP Echo Requests to the virtual router instance IP addresses will be silently discarded. The no form of this command restores the default operation of discarding all ICMP Echo Request messages destined to the non-owner virtual router instance IP addresses.
Default	no ping-reply

policy

Syntax	policy <i>vrrp-policy-id</i> no policy
Context	config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp
Description	This command associates a VRRP priority control policy with the virtual router instance (non-owner context only).
Parameters	<i>vrrp-policy-id</i> — Specifies a VRRP priority control policy.
Values	1 to 9999

preempt

Syntax	[no] preempt
Context	config>service>vprn>if config>service>vprn>if>ipv6>vrrp
Description	The preempt mode value controls whether a specific backup virtual router preempts a lower priority master. When preempt is enabled, the virtual router instance overrides any non-owner master with an “in use” message priority value less than the virtual router instance in-use priority value. If preempt is disabled, the virtual router only becomes master if the master down timer expires before a VRRP advertisement message is received from another virtual router.

The IP address owner will always become master when available. Preempt mode cannot be disabled on the owner virtual router.

The default value for preempt mode is enabled.

Default preempt

priority

Syntax **priority** *priority*
no priority

Context config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp

Description The priority command provides the ability to configure a specific priority value to the virtual router instance. In conjunction with an optional policy command, the base-priority is used to derive the in-use priority of the virtual router instance.

The priority command is only available in the non-owner **vrrp** *virtual-router-id* nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed. For non-owner virtual router instances, if the priority command is not executed, the base-priority will be set to 100.

The **no** form of this command restores the default value of 100 to base-priority.

Parameters *base-priority* — The base-priority parameter configures the base priority used by the virtual router instance. If a VRRP priority control policy is not also defined, the base-priority will be the in-use priority for the virtual router instance.

Values 1 to 254

Default 100

ssh-reply

Syntax [**no**] **ssh-reply**

Context config>service>vprn>if>vrrp

Description This command enables the non-owner master to reply to SSH Requests directed at the virtual router instance's IP addresses. The SSH request can be received on any routed interface. SSH must not have been disabled at the management security level (either on the parental IP interface or based on the SSH source host address). Proper login and CLI command authentication is still enforced.

When ssh-reply is not enabled, SSH packets to non-owner master virtual IP addresses are silently discarded. Non-owner backup virtual routers never respond to SSH regardless of the ssh-reply configuration.

The `ssh-reply` command is only available in non-owner **vrrp** *virtual-router-id* nodal context. If the `ssh-reply` command is not executed, SSH packets to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all SSH packets destined to the non-owner virtual router instance IP addresses.

Default no ssh-reply

standby-forwarding

Syntax **[no] standby-forwarding**

Context config>service>vprn>if>vrrp
 config>service>vprn>if>ipv6>vrrp

Description This command allows the forwarding of packets by a standby router.

The **no** form of the command specifies that a standby router should not forward traffic sent to virtual router's MAC address. However, the standby router should forward traffic sent to the standby router's real MAC address.

Default no standby-forwarding

telnet-reply

Syntax **[no] telnet-reply**

Context config>service>vprn>if>vrrp
 config>service>vprn>if>ipv6>vrrp

Description This command enables the non-owner master to reply to TCP port 23 Telnet Requests directed at the virtual router instance's IP addresses. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Proper login and CLI command authentication is still enforced.

When `telnet-reply` is not enabled, TCP port 23 Telnet packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to Telnet Requests regardless of the `telnet-reply` configuration.

The `telnet-reply` command is only available in non-owner **VRRP** nodal context. If the `telnet-reply` command is not executed, Telnet packets to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all Telnet packets destined to the non-owner virtual router instance IP addresses.

Default no telnet-reply

traceroute-reply

Syntax [no] traceroute-reply

Context config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp

Description This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.

When this command is enabled, a non-owner master can reply to traceroute requests directed to the virtual router instance IP addresses.

A non-owner backup virtual router never responds to such traceroute requests regardless of the **trace-route-reply** status.

Default no traceroute-reply

3.8.2.24 Interface SAP Commands

sap

Syntax sap sap-id [create]
no sap sap-id

Context config>service>vprn>if
config>service>vprn>sub-if>grp-if
config>service>vprn>if>sap

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object. Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface port-type port-id mode access** command. Channelized TDM ports are always access ports.

If a port is shut down, all SAPs on that port become operationally down. When a service is shut down, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP will also be deleted. The **no** form of the command causes the ptp-h-assist to be disabled.

Default No SAPs are defined.

Special Cases **VPRN** — A VPRN SAP must be defined on an Ethernet interface.

sap ipsec-id.private | public:tag — This parameter associates an IPsec group SAP with this interface. This is the public side for an IPsec tunnel. Tunnels referencing this IPsec group in the private side may be created if their local IP is in the subnet of the interface subnet and the routing context specified matches with the one of the interface.

This context will provide a SAP to the tunnel. The operator may associate an ingress and egress QoS policies as well as filters and virtual scheduling contexts. Internally this creates an Ethernet SAP that will be used to send and receive encrypted traffic to and from the MDA. Multiple tunnels can be associated with this SAP. The “tag” will be a dot1q value. The operator may see it as an identifier. The range is limited to 1 to 4094.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

port-id — Specifies the physical port ID.

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 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels the port ID must include the channel ID. A period “.” separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

<i>port-id</i>	<i>slot/mda/port</i> [<i>channel</i>]	
	eth-sat-id	esat-id/slot/port
		esat keyword
		id 1 to 20
	pxc-id	pxc-id.sub-port
		pxc keyword
		id 1 to 64
		sub-port a, b

create — Creates a SAP instance.

split-horizon-group *group-name* — Specifies the name of the split horizon group to which the SAP belongs.

aarp

Syntax	aarp <i>aarpId</i> type <i>type</i> no aarp
Context	config>service>vprn>if>sap config>service>vprn>if>spoke-sdp
Description	This command associates an AARP instance with a multi-homed SAP or spoke SDP. This instance uses the same AARP ID in the same node or in a peer node (pre-configured) to provide traffic flow and packet asymmetry removal for a multi-homed SAP or spoke SDP. The type specifies the role of this service point in the AARP: either, primary (dual-homed) or secondary (dual-homed-secondary). The AA service attributes (app-profile and transit-policy) of the primary are inherited by the secondary endpoints. All endpoints within an AARP must be of the same type (SAP or spoke), and all endpoints with an AARP must be within the same service. The no form of the command removes the association between an AARP instance and a multi-homed SAP or spoke SDP.
Default	no aarp
Parameters	aarpId — Specifies the AARP instance associated with this SAP. If not configured, no AARP instance is associated with this SAP. Values 1 to 65535 type — Specifies the role of the SAP referenced by the AARP instance. Values dual-homed — The primary dual-homed AA subscriber side service-point of an AARP instance; only supported for Epipe, IES, and VPRN SAP and spoke SDP. dual-homed-secondary — One of the secondary dual-homed AA subscriber side service-points of an AARP instance; only supported for Epipe, IES, and VPRN SAP and spoke SDP.

transit-policy

Syntax	transit-policy { ip <i>ip-aasub-policy-id</i> prefix <i>prefix-aasub-policy-id</i> } no transit-policy
Context	config>service>vprn>if>sap> config>service>vprn>if>spoke-sdp>

Description	This command associates an AA transit policy to the service. The transit IP policy must be defined prior to associating the policy with a SAP in the config>application- assurance>group>policy>transit-ip-policy context. Transit AA subscribers are managed by the system through this service policy, which determines how transit subs are created and removed for that service. The no form of the command removes the association of the policy to the service.
Default	no transit-policy
Parameters	<i>ip-aasub-policy-id</i> — Specifies an integer identifying an IP transit IP profile entry. Values 1 to 65535 <i>prefix-aasub-policy-id</i> — Specifies an integer identifying a prefix transit profile entry. Values 1 to 65535

accounting-policy

Syntax	accounting-policy <i>acct-policy-id</i> no accounting-policy
Context	config>service>vprn>if>sap config>service>vprn>if>spoke-sdp
Description	This command creates the accounting policy context that can be applied to an interface SAP or interface SAP spoke SDP. An accounting policy must be defined before it can be associated with a SAP. If the <i>policy-id</i> 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>accounting-policy context. The no form of this command removes the accounting policy association from the SAP, and the accounting policy reverts to the default.
Default	no accounting-policy
Parameters	<i>acct-policy-id</i> — The accounting <i>policy-id</i> as configured in the config>log>accounting-policy context. Values 1 to 99

app-profile

Syntax	app-profile <i>app-profile-name</i> no app-profile

Context	config>service>vprn>if>spoke-sdp
Description	This command configures the application profile name.
Parameters	<i>app-profile-name</i> — Specifies the application profile name.
Values	32 chars max

bfd-enable

Syntax	bdf-enable no bfd-enable
Context	config>service>vprn>if>spoke-sdp
Description	This command enables VCCV BFD on the PW associated with the VLL, BGP VPWS, or VPLS service. The parameters for the BFD session are derived from the named BFD template, which must have been first configured using the bfd-template command.

bfd-template

Syntax	bfd-template <i>name</i> no bfd-template
Context	config>service>vprn>if>spoke-sdp
Description	This command configures a named BFD template to be used by VCCV BFD on PWs belonging to the VLL, BGP VPWS, or VPLS service. The template specifies parameters, such as the minimum transmit and receive control packet timer intervals, to be used by the BFD session. Template parameters are configured under the config>router>bfd context.
Default	no bfd-template
Parameters	<i>name</i> — A text string name for the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

collect-stats

Syntax	[no] collect-stats
Context	config>service>vprn>if>sap config>service>vprn>if>spoke-sdp
Description	This command enables accounting and statistical data collection for either an interface SAP or interface SAP spoke SDP, or network port. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

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

cpu-protection

Syntax **cpu-protection** *policy-id* [*mac-monitoring*] | [*eth-cfm-monitoring* [*aggregate*] [*car*]]
no cpu-protection

Context config>service>vprn>sub-if>grp-if>sap

Description This command assigns an existing CPU protection policy to the associated group interface. The CPU protection policies are configured in the **config>sys>security>cpu-protection>policy** *cpu-protection-policy-id* context.

If no CPU-Protection policy is assigned to a group interface SAP, then the default policy is used to limit the overall-rate. The default policy is policy number 254 for access interfaces and 255 for network interfaces.

The no form of the command removes the association of the CPU protection policy from the associated interface and reverts to the default policy values.

Default cpu-protection 254 (for access interfaces); cpu-protection 255 (for network interfaces)

The configuration of no cpu-protection returns the interface/SAP to the default policies as shown above.

Parameters *policy-id* — Specifies an existing CPU protection policy.

Values 1 to 255

mac-monitoring — Enables MAC monitoring.

eth-cfm-monitoring — Enables Ethernet Connectivity Fault Management monitoring.

aggregate — Applies the rate limit to the sum of the per peer packet rates.

car — (Committed Access Rate) *causes* Eth-CFM packets to be ignored when enforcing the overall-rate.

dist-cpu-protection

Syntax **dist-cpu-protection** *policy-name*
no dist-cpu-protection

Context config>service>vprn>sub-if>grp-if>sap
config>service>vprn>if>sap

Description	This command assigns a Distributed CPU Protection (DCP) policy to the SAP. Only a valid created DCP policy can be assigned to a SAP or a network interface (This rule does not apply to templates such as an msap-policy).
Default	If no dist-cpu-protection policy is assigned to an SAP policy, then the default access DCP policy (default-access-policy) is used. If no DCP functionality is required on the SAP policy, then an empty DCP policy can be created and explicitly assigned to the SAP policy.
Parameters	<i>policy-name</i> — Specifies the name of the DCP policy up to 32 characters in length

default-host

Syntax	default-host <i>ip-address/mask</i> next-hop <i>next-hop-ip</i> no default-host <i>ip-address/mask</i>
Context	config>service>vprn>sub-if>grp-if>sap
Description	This command configures the default-host to be used. More than one default-host can be configured per SAP. The no form of the command removes the values from the configuration.
Parameters	<i>ip-address/mask</i> — Assigns an IP address/IP subnet format to the interface. next-hop <i>next-hop-ip</i> — Assigns the next hop IP address.

source

Syntax	source <i>ip-address</i> no source
Context	config>service>ies>if>sap>ip-tunnel
Description	This command configures the source IPv4 or IPv6 address to use for an IP tunnel. This configuration applies to the outer IP header of the encapsulated packets. The IPv4 or IPv6 address must belong to the one of the IP subnets associated with the public SAP interface of the tunnel-group. The source address, remote-ip address and backup-remote-ip address of a tunnel must all belong to the same address family (IPv4 or IPv6). When the source address contains an IPv6 address it must be a global unicast address.
Default	no source
Parameters	<i>ip-address</i> — An IPv4 address or an IPv6 address.

remote-ip

Syntax	remote-ip <i>ip-address</i> no remote-ip
Context	config>service>ies>if>sap>ip-tunnel
Description	This command configures the primary destination IPv4 or IPv6 address to use for an IP tunnel. This configuration applies to the outer IP header of the encapsulated packets. The source address, remote-ip address and backup-remote-ip address of a tunnel must all belong to the same address family (IPv4 or IPv6). When the remote-ip address contains an IPv6 address it must be a global unicast address.
Default	no remote-ip
Parameters	<i>ip-address</i> — An IPv4 address or an IPv6 address.

backup-remote-ip

Syntax	backup-remote-ip <i>ip-address</i> no remote-ip
Context	config>service>ies>if>sap>ip-tunnel
Description	This command configures the alternate destination IPv4 or IPv6 address to use for an IP tunnel. This destination address is used only if the primary destination configured with the remote-ip command is unreachable in the delivery service. The source address, remote-ip address and backup-remote-ip address of a tunnel must all belong to the same address family (IPv4 or IPv6). When the backup-remote-ip address contains an IPv6 address it must be a global unicast address.
Default	no backup-remote-ip
Parameters	<i>ip-address</i> — An IPv4 address or an IPv6 address.

3.8.2.25 Routed VPLS Commands

vpls

Syntax	vpls <i>service-name</i>
Context	config>service config>service>vprn>if
Description	The vpls command, within the IP interface context, is used to bind the IP interface to the specified service name.

The system does not attempt to resolve the service name provided until the IP interface is placed into the administratively up state (**no shutdown**). Once the IP interface is administratively up, the system will scan the available VPLS services that have the **allow-ip-int-bind** flag set for a VPLS service associated with the name. If the service name is bound to the service name when the IP interface is already in the administratively up state, the system will immediately attempt to resolve the given name.

If a VPLS service is found associated with the name and with the **allow-ip-int-bind** flag set, the IP interface will be attached to the VPLS service allowing routing to and from the service virtual ports once the IP interface is operational.

A VPLS service associated with the specified name that does not have the **allow-ip-int-bind** 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-bind** flag is set at the time the name is applied, the VPLS service will be 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-bind** flag set, the system will 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-bind** flag set will be attached to the VPLS service. Only one IP interface is allowed to attach to a VPLS service context. No error is generated for the remaining non-attached IP interfaces using the service name.

Once an IP interface is attached to a VPLS service, the name associated with the service cannot be removed or changed until the IP interface name binding is removed. Also, the **allow-ip-int-bind** 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.

VPRN Hardware Dependency

When a service name is bound to a VPRN IP interface, all SAPs associated with the VPRN service must be on hardware based on the FlexPath forwarding plane. Currently, these include the IOM3-XP, the various IMM modules and the SR7710c12. If any SAPs are associated with the wrong hardware type, the service name binding to the VPRN IP interface fails. Once an IP interface within the VPRN service is bound to a service name, attempting to create a SAP on excluded hardware fails.

IP Interface MTU and Fragmentation

A 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 and CRC) that may be transmitted out a port. The VPLS itself has a service level MTU that defines the largest packet supported by the service. This MTU does not include the local encapsulation overhead for each port (QinQ, Dot1Q, TopQ or SDP service delineation fields and headers) but does include the remainder of the packet. As virtual ports are created in the system, the virtual port cannot become operational unless the configured port MTU minus the virtual port service delineation overhead is greater than or equal to the configured VPLS service MTU. Thus, an operational virtual port 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.

IP interfaces have a configurable up MTU that defines the largest packet that may egress the IP interface without being fragmented. This MTU encompasses the IP portion of the packet and does not include any of the egress DLC header or CRC. This MTU does not affect the size of the largest ingress packet on the IP interface. If the egress IP portion of the packet is larger than the IP interface MTU and the IP header do not fragment flag is not set, the packet is fragmented into smaller packets that will not exceed the configured MTU size. If the do not fragment bit is set, the packet is silently discarded at egress when it exceeds the IP MTU.

When the IP interface is bound to a VPLS service, the IP MTU must be at least 18 bytes less than the VPLS service MTU. This allows for the addition of the minimal Ethernet encapsulation overhead; 6 bytes for the DA, 6 bytes for the SA, 2 bytes for the Etype and 4 bytes for the trailing CRC. Any remaining egress virtual port overhead (Dot1P, Dot1Q, QinQ, TopQ or SDP) required above the minimum is known to be less than the egress ports MTU since the virtual port would not be operational otherwise.

If the IP interface IP MTU value is too large based on the VPLS service MTU, the IP interface will enter the operationally down state until either the IP MTU is adequately lowered or the VPLS service MTU is sufficiently increased.

The **no** form of the command on the IP interface is used to remove the service name binding from the IP interface. If the service name has been resolved to a VPLS service context and the IP interface has been attached to the VPLS service, the IP interface will also be detached from the VPLS service.

Parameters *service-name* — The service-name parameter is required when using the IP interface vpls command and specifies the service name that the system will attempt to resolve to an **allow-ip-int-bind** 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.

egress

Syntax **egress**

Context config>service>vprn>if>vpls

Description The egress node under the vpls binding is used to define the optional sap-egress QoS policy that will be used for reclassifying the egress forwarding class or profile for routed packets associated with the IP interface on the attached VPLS service context.

ingress

Syntax **ingress**

Context config>service>vprn>if>vpls

Description The ingress node in this context under the vpls binding is used to define the routed IPv4 and IPv6 optional filter overrides.

v4-routed-override-filter

Syntax **v4-routed-override-filter** *ipv4-filter-id*
no v4-routed-override-filter

Context config>service>vprn>if>vpls>egress

Description This command configures an IPv4 filter ID that is applied to packets egressing the VPRN R-VPLS interface. The filter overrides the existing egress IPv4 filter applied to VPLS service endpoints such as SAPs or SDPs, if configured.

 The **no** form of the command removes the IPv4 routed override filter from the egress VPRN R-VPLS interface. When removed, egress IPv4 packets will use the IPv4 egress filter applied to VPLS endpoint, if configured.

Parameters *ip-filter-id* — Specifies the IP filter ID. This parameter is required when executing the **v4-routed-override-filter** command. The specified filter ID must exist as an IPv4 filter within the system or the override command fails.

v4-routed-override-filter

Syntax	v4-routed-override-filter <i>ipv4-filter-id</i> no v4-routed-override-filter
Context	config>service>vprn>if>vpls>ingress
Description	This command configures an IPv4 filter ID that is applied to all ingress packets entering the VPLS service. The filter overrides any existing ingress IPv4 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 IPv4 routed packet's will use the any existing ingress IPv4 filter on the VPLS virtual port. The no form of the command removes the IPv4 routed override filter from the ingress IP interface. When removed, the IPv4 ingress routed packets within a VPLS service attached to the IP interface will use the IPv4 ingress filter applied to the packets virtual port, when defined.
Parameters	<i>ip-filter-id</i> — Specifies the IP filter ID. This parameter is required when executing the v4-routed-override-filter command. The specified filter ID must exist as an IPv4 filter within the system or the override command fails.

v6-routed-override-filter

Syntax	v6-routed-override-filter <i>ipv6-filter-id</i> no v6-routed-override-filter
Context	config>service>vprn>if>vpls>egress
Description	This command configures an IPv6 filter ID that is applied to packets egressing the VPRN R-VPLS interface. The filter overrides existing egress IPv6 filter applied to VPLS service endpoints such as SAPs or SDPs, if configured. The no form of the command removes the IPv4 routed override filter from the egress VPRN R-VPLS interface. When removed, egress IPv6 packets will use the IPv6 egress filter applied to the VPLS endpoint, if configured.
Parameters	<i>ipv6-filter-id</i> — Specifies the IPv6 filter ID. This parameter is required when executing the v6-routed-override-filter command. The specified filter ID must exist as an IPv6 filter within the system or the override command fails.

v6-routed-override-filter

Syntax	v6-routed-override-filter <i>ipv6-filter-id</i> no v6-routed-override-filter
Context	config>service>vprn>if>vpls>ingress

Description	This command configures an IPv6 filter ID that is applied to all ingress packets entering the VPLS service. The filter overrides any existing ingress IPv6 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 IPv6 routed packets use the any existing ingress IPv6 filter on the VPLS virtual port. The no form of the command removes the IPv6 routed override filter from the ingress IP interface. When removed, the IPv6 ingress routed packets within a VPLS service attached to the IP interface uses the IPv6 ingress filter applied to the packet's virtual port, when defined.
Parameters	<i>ipv6-filter-id</i> — Specifies the IPv6 filter ID. This parameter is required when executing the v6-routed-override-filter command. The specified filter ID must exist as an IPv6 filter within the system or the override command fails.

reclassify-using-qos

Syntax	reclassify-using-qos <i>policy-id</i> no reclassify-using-qos
Context	config>service>vprn>if>vpls>egress
Description	This command specifies a SAP egress QoS policy that is used to reclassify the forwarding class and profile of egress routed packets on the VPLS service. When routed packets associated with the IP interface egress a VPLS SAP, the reclassification rules within the sap-egress QoS policy applied to the SAP are always ignored (even when reclassify-using-qos is not defined). Any queues or policers defined within the specified QoS policy are ignored and are not created on the VPLS egress SAPs. Instead, the routed packets continue to use the forwarding class mappings, queues and policers from the SAP egress QoS policy applied to the egress VPLS SAP. While the specified SAP egress policy ID is applied to an IP interface it cannot be deleted from the system. The no form of the command removes the SAP egress QoS policy used for reclassification from the egress IP interface. When removed, IP routed packets is not reclassified on the egress SAPs of the VPLS service attached to the IP interface.
Parameters	<i>policy-id</i> — Specifies the SAP egress QoS policy ID This parameter is required when executing the reclassify-using-qos command. The specified SAP egress QoS ID must exist within the system or the command fails.

3.8.2.26 Network Ingress Commands

network

Syntax	network
Context	config>service>vprn
Description	This command enters the context to configure network parameters for the VPRN service.

ingress

Syntax	ingress
Context	config>service>vprn>network
Description	This command enters the context to configure network ingress parameters for the VPRN service.

qos

Syntax	qos <i>network-policy-id</i> fp-redirect-group <i>queue-group-name</i> <i>instance</i> <i>instance-id</i> no qos
Context	config>service>vprn>network>ingress
Description	This command is used to redirect unicast packets arriving on an automatically (using the auto-bind-tunnel command) or manually configured (using a spoke-sdp command, but not the spoke-sdp command under the VPRN IP interface) binding in a VPRN to a policer in an ingress forwarding plane queue-group for the purpose of rate-limiting.

For the policer to be used, the following must be true:

- Step 1.** The configured queue group template name must be applied to the forwarding plane on which the ingress traffic arrives using the instance id specified.
- Step 2.** The policer referenced in the FC-to-policer mappings in the ingress context of a network QoS policy must be present in the specified queue group template.

The command fails if the queue group template name does not exist or if the policer specified in the network QoS policy does not exist in the queue group template. If the queue group template name with the specified instance is not applied to the forwarding plane on which the VPRN binding unicast traffic arrives then this traffic uses the ingress network queues related to the network interface, however, the ingress classification is still based on the applied network QoS policy.

The unicast traffic can be redirected to a policer under the forwarding class **fp-redirect-group** command in the ingress section of a network QoS policy; any **fp-redirect-group multicast-policer**, **broadcast-policer** or **unknown-policer** commands are ignored for this traffic. Multicast traffic would use the ingress network queues or queue group related to the network interface.

Ingress classification is based on the configuration of the ingress section of the specified network QoS policy, noting that the dot1p and exp classification is based on the outer Ethernet header and MPLS label whereas the DSCP applies to the outer IP header if the tunnel encapsulation is GRE, or the DSCP in the first IP header in the payload if **ler-use-dscp** is enabled in the ingress section of the referenced network QoS policy.

When this command is applied, it overrides the QoS applied to the related network interfaces for unicast traffic arriving on bindings in that VPRN.

The **no** version of this command removes the redirection of VPRN binding traffic to the queue-group policers.

Parameters	<i>network-policy-id</i> — Specifies the network policy identification. The value uniquely identifies the policy on the system.
	Values 1 to 65535
	fp-redirect-group <i>queue-group-name</i> — Specifies the name of the queue group template up to 32 characters in length.
	instance <i>instance-id</i> — Specifies the identification of a specific instance of the queue-group.
	Values 1 to 65535

urpf-check

Syntax	urpf-check no urpf-check
Context	config>service>vprn>network>ingress
Description	This command enables the unicast RPF (uRPF) check of network ingress traffic to include traffic associated with the VPRN if the incoming network interface is configured with the urpf-selected-vprns command If the command is not configured, then traffic associated with this VPRN that arrives on a network interface with urpf-selected-vprns configured bypasses the uRPF checking options specified for that network interface.
Default	no urpf-check

3.8.2.27 SAP Subscriber Management Commands

sub-sla-mgmt

Syntax	[no] sub-sla-mgmt
Context	config>service>vprn>sub-if>grp-if>sap
Description	This command enters the context to configure subscriber management parameters for this SAP.
Default	no sub-sla-mgmt

def-sla-profile

Syntax	def-sla-profile default-sla-profile-name no def-sla-profile
Context	config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
Description	This command specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the config>subscr-mgmt>sla-profile context. An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts. SLA profiles are maintained in two locations, the subscriber identification policy and the subscriber profile templates. After a subscriber host is associated with an SLA profile name, either the subscriber identification policy used to identify the subscriber or the subscriber profile associated with the subscriber host must contain an SLA profile with that name. If both the subscriber identification policy and the subscriber profile contain the SLA profile name, the SLA profile in the subscriber profile is used. The no form of the command removes the default SLA profile from the SAP configuration.
Default	no def-sla-profile
Parameters	<i>default-sla-profile-name</i> — Specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the config>subscr-mgmt>sla-profile context.

def-sub-profile

Syntax	def-sub-profile default-subscriber-profile-name
Context	config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt

Description This command specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sub-profile** context.

A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscriber using the subscriber profile. Subscriber profiles also allow for specific SLA profile definitions when the default definitions from the subscriber identification policy must be overridden.

The **no** form of the command removes the default SLA profile from the SAP configuration.

Parameters *default-sub-profile* — Specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sub-profile** context.

multi-sub-sap

Syntax **multi-sub-sap** [*number-of-sub*]
no multi-sub-sap

Context config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt

Description This command configures the maximum number of subscribers for this SAP. It is used in conjunction with the profiled-traffic-only command on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues.

The **no** form of this command returns the default value.

Default multi-sub-sap 1

Parameters *number-of-sub* — Specifies the maximum number of subscribers for this SAP.

Values 2 to 8000

single-sub-parameters

Syntax **single-sub-parameters**

Context config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt

Description This command enters the context to configure single subscriber parameters for this SAP.

non-sub-traffic

Syntax **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-*

*ident-string]***no non-sub-traffic**

Context	config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt>single-sub
Description	This command configures non-subscriber traffic profiles. It is used in conjunction with the <code>profiled-traffic-only</code> command on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues. The no form of the command removes the profiles and disables the feature.
Parameters	sub-profile <i>sub-profile-name</i> — Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the config>subscr-mgmt>sub-profile context. sla-profile <i>sla-profile-name</i> — Specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the config>subscr-mgmt>sla-profile context. subscriber <i>sub-ident-string</i> — Specifies an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the config>subscr-mgmt>sub-ident-policy context. The subscriber information is used by the VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers. For VPRN SAPs with arp-reply-agent enabled with the optional <i>sub-ident</i> parameter, the static subscriber host's <i>sub-ident-string</i> is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPRN destinations. If the static subscriber host's <i>sub-ident</i> string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP. If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP. If <i>sub-ident</i> is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP. ARP requests are never forwarded back to the same SAP or within the receiving SAP's split horizon group.

profiled-traffic-only

Syntax	[no] profiled-traffic-only
Context	config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt>single-sub

Description This command enables profiled traffic only for this SAP. The profiled traffic refers to single subscriber traffic on a dedicated SAP (in the VLAN-per-subscriber model). When enabled, subscriber queues are instantiated through the QOS policy defined in the sla-profile and the associated SAP queues are deleted. This can increase subscriber scaling by reducing the number of queues instantiated per subscriber (in the VLAN-per-subscriber model). In order for this to be achieved, any configured multi-sub-sap limit must be removed (leaving the default of 1).

The **no** form of the command disables the command.

sub-ident-policy

Syntax **sub-ident-policy** *sub-ident-policy-name*

Context config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt

Description This command associates a subscriber identification policy to this SAP. The subscriber identification policy must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sub-ident-policy** context.

Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.

For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet string. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.

When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.

The **no** form of the command removes the default subscriber identification policy from the SAP configuration.

Default no sub-ident-policy

Parameters *sub-ident-policy-name* — Specifies a subscriber identification policy for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sub-ident-policy** context.

srrp

Syntax [**no**] **srrp** *srrp-id*

Context config>service>vprn>sub-if>grp-if

Description	This command creates an SRRP instance on a group IP interface. An SRRP instance manages all subscriber subnets within the group interfaces subscriber IP interface or other subscriber IP interfaces that are associated through a wholesale/retail relationship. Only one unique SRRP instance can be configured per group interface. The no form of the command removes an SRRP instance from a group IP interface. Once removed, the group interface ignores ARP requests for the SRRP gateway IP addresses that may exist on subscriber subnets associated with the group IP interface. Then the group interface stops routing using the redundant IP interface associated with the group IP interface and stops routing with the SRRP gateway MAC address. Ingress packets destined to the SRRP gateway MAC is also silently discarded. This is the same behavior as a group IP interface that is disabled (shutdown).
Default	no srrp
Parameters	<i>srrp-id</i> — Specifies a 32 bit instance ID that must be unique to the system. The instance ID must also match the instance ID used by the remote router that is participating in the same SRRP context. SRRP is intended to perform a function similar to VRRP where adjacent IP hosts within local subnets use a default gateway to access IP hosts on other subnets. Values 1 to 4294967295

gw-mac

Syntax	gw-mac <i>mac-address</i> no gw-mac
Context	config>service>vprn>sub-if>grp-if>srrp
Description	This command overrides the default SRRP gateway MAC address used by the SRRP instance. Unless specified, the system uses the same base MAC address for all SRRP instances with the last octet overridden by the lower 8 bits of the SRRP instance ID. The same SRRP gateway MAC address should be in-use by both the local and remote routers participating in the same SRRP context. One reason to change the default SRRP gateway MAC address is if two SRRP instances sharing the same broadcast domain are using the same SRRP gateway MAC. The system uses the SRRP instance ID to separate the SRRP messages (by ignoring the messages that does not match the local instance ID), but a unique SRRP gateway MAC is essential to separate the routed packets for each gateway IP address. The no form of the command removes the explicit SRRP gateway MAC address from the SRRP instance. The SRRP gateway MAC address can only be changed or removed when the SRRP instance is shutdown.

Parameters	<i>mac-address</i> — Specifies a MAC address that is used to override the default SRRP base MAC address.
Values	Any MAC address except all zeros, broadcast or multicast addresses. The offset is expressed in normal Ethernet MAC address notation. The defined gw-mac cannot be 00:00:00:00:00:00, ff:ff:ff:ff:ff:ff or any multicast address.
Default	If not specified, the system uses the default SRRP gateway MAC address with the last octet set to the 8 least significant bits of the SRRP instance ID.

keep-alive-interval

Syntax	keep-alive-interval <i>interval</i> no keep-alive-interval
Context	config>service>vprn>sub-if>grp-if>srrp
Description	This command defines the interval between SRRP advertisement messages sent when operating in the master state. The interval is also the basis for setting the master-down timer used to determine when the master is no longer sending. The system uses three times the keep-alive interval to set the timer. Every time an SRRP advertisement is seen that is better than the local priority, the timer is reset. If the timer expires, the SRRP instance assumes that a master does not exist and initiates the attempt to become master. When in backup state, the SRRP instance takes the keep-alive interval of the master as represented in the masters SRRP advertisement message. Once in master state, the SRRP instance uses its own configured keep-alive interval. The keep-alive-interval may be changed at anytime, but has no effect until the SRRP instance is in the master state. The no form of the command restores the default interval.
Parameters	<i>interval</i> — Specifies the interval, in ms, between SRRP advertisement messages sent when operating in the master state.
Values	1 to 100
Default	10 ms

message-path

Syntax	message-path <i>sap-id</i> no message-path
Context	config>service>vprn>sub-if>grp-if>srrp

Description This command defines a specific SAP for SRRP in-band messaging. A message-path SAP must be defined prior to activating the SRRP instance. The defined SAP must exist on the SRRP instances group IP interface for the command to succeed and cannot currently be associated with any dynamic or static subscriber hosts. Once a group IP interface SAP has been defined as the transmission path for SRRP Advertisement messages, it cannot be administratively shut down, will not support static or dynamic subscriber hosts and cannot be removed from the group IP interface.

The SRRP instance message-path command may be executed at anytime on the SRRP instance. Changing the message SAP fails if a dynamic or static subscriber host is associated with the new SAP. Once successfully changed, the SRRP instance immediately disables anti-spoofing on the SAP and start sending SRRP Advertisement messages if the SRRP instance is activated.

Changing the current SRRP message SAP on an active pair of routers should be done in the following manner:

1. Shut down the backup SRRP instance.
2. Change the message SAP on the **shutdown** node.
3. Change the message SAP on the active master node.
4. Re-activate the shut down SRRP instance.

Shutting down the backup SRRP instance prevents the SRRP instances from becoming master due to temporarily using differing message path SAPs.

If an MCS peering is operational between the redundant nodes and the SRRP instance has been associated with the peering, the designated message path SAP is sent from each member.

The **no** form of the command can only be executed when the SRRP instance is shut down. Executing no message-path allows the existing SAP to be used for subscriber management functions. A new message-path SAP must be defined prior to activating the SRRP instance.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

policy

Syntax **[no] policy** *vrrp-policy-id*

Context config>service>vprn>sub-if>grp-if>srrp

Description This command associates one or more VRRP policies with the SRRP instance. A VRRP policy is a collection of connectivity and verification tests used to manipulate the in-use priorities of VRRP and SRRP instances. A VRRP policy can test the link state of ports, ping IP hosts, discover the existence of routes in the routing table or the ability to reach Layer 2 hosts. When one or more of these tests fail, the VRRP policy has the option of decrementing or setting an explicit value for the in-use priority of an SRRP instance.

More than one VRRP policy may be associated with an SRRP instance. When more than one VRRP policy is associated with an SRRP instance the delta decrement of the in-use priority is cumulative unless one or more test fail that have explicit priority values. When one or more explicit tests fail, the lowest priority value event takes effect for the SRRP instance. When the highest delta-in-use-limit is used to manage the lowest delta derived in-use priority for the SRRP instance.

VRRP policy associations may be added and removed at anytime. A maximum of two VRRP policies can be associated with a single SRRP instance.

The **no** form of the command removes the association with vrrp-policy-id from the SRRP instance.

Parameters

vrrp-policy-id — Specifies one or more VRRP policies with the SRRP instance.

Values1 to 9999

priority

Syntaxpriority priority
no priority

Contextconfig>service>vprn>sub-if>grp-if>srrp

DescriptionThis command overrides the default base priority for the SRRP instance. The SRRP instance priority is advertised by the SRRP instance to its neighbor router and is compared to the priority received from the neighbor router. The router with the best (highest) priority enters the master state while the other router enters the backup state. If the priority of each router is the same, the router with the lowest source IP address in the SRRP advertisement message assumes the master state.

The base priority of an SRRP instance can be managed by VRRP policies. A VRRP policy defines a set of connectivity or verification tests which, when they fail, may lower an SRRP instances base priority (creating an in-use priority for the instance). Every time an SRRP instances in-use priority changes when in master state, it sends an SRRP advertisement message with the new priority. If the dynamic priority drops to zero or receives an SRRP Advertisement message with a better priority, the SRRP instance transitions to the *becoming backup* state.

When the priority command is not specified, or the no priority command is executed, the system uses a default base priority of 100. The priority command may be executed at anytime.

The **no** form of the command restores the default base priority to the SRRP instance. If a VRRP policy is associated with the SRRP instance, it uses the default base priority as the basis for any modifications to the SRRP instances in-use priority.

Parameters	<i>priority</i> — Specifies a base priority for the SRRP instance to override the default.
Values	1 to 254
Default	100

send-fib-population-packets

Syntax	send-fib-population-packets (all outer-tag-only) no send-fib-population-packets
Context	config>service>vprn>sub-if>grp-if>srrp
Description	This command sends FIB population packets. The no form of the command disables sending FIB population packets.
Default	send-fib-population-packets all
Parameters	<i>all</i> — Sends FIB population packets to all VLANs. <i>outer-tag-only</i> — Sends FIB population packets to only outer VLAN tags.

3.8.2.28 BGP Commands

bgp

Syntax	[no] bgp
Context	service>vprn
Description	This command enables the BGP protocol with the VPRN service. The no form of the command disables the BGP protocol from the given VPRN service.
Default	no bgp

bgp-shared-queue

Syntax	bgp-shared-queue [<i>cir rate</i>] [<i>pir rate</i>] no bgp-shared-queue
Context	config>service>vprn

Description	This command enables all BGP peers within a VPRN instance to share a single CPM queue. This command takes effect on new BGP connections established; already established BGP peers continue to use their own CPM queue. Any changes to PIR/CIR of the shared queue takes effect only after BGP connections are re-established.
Parameters	cir rate — Specifies the CIR rate for the shared queue. pir rate — Specifies the PIR rate for the shared queue.

advertise-inactive

Syntax	[no] advertise-inactive
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command enables or disables the advertising of inactive BGP routers to other BGP peers. By default, BGP only advertises BGP routes to other BGP peers if a given BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the most preferred route within the system for a given destination. When the BGP advertise-inactive command is configured so that it applies to a BGP session it has the following effect on the IPv4, IPv6, mcast-ipv4, mcast-ipv6, label-IPv4 and label-IPv6 routes advertised to that peer: • If the active route for the IP prefix is a BGP route then that route is advertised. • If the active route for the IP prefix is a non-BGP route and there is at least one valid but inactive BGP route for the same destination then the best of the inactive and valid BGP routes is advertised unless the non-BGP active route is matched and accepted by an export policy applied to the session. • If the active route for the IP prefix is a non-BGP route and there are no (valid) BGP routes for the same destination then no route is advertised for the prefix unless the non-BGP active route is matched and accepted by an export policy applied to the session.
Default	no advertise-inactive

aggregator-id-zero

Syntax	[no] aggregator-id-zero
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor

Description	This command is used to set the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes that contain different AS paths. When BGP is aggregating routes, it adds the aggregator path attribute to the BGP update messages. By default, BGP adds the AS number and router ID to the aggregator path attribute. When this command is enabled, BGP adds the router ID to the aggregator path attribute. This command is used at the group level to revert to the value defined under the global level, while this command is used at the neighbor level to revert to the value defined under the group level. The no form of the command used at the global level reverts to default where BGP adds the AS number and router ID to the aggregator path attribute. The no form of the command used at the group level reverts to the value defined at the group level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	no aggregator-id-zero — BGP adds the AS number and router ID to the aggregator path attribute.

always-compare-med

Syntax	always-compare-med {zero infinity} no always-compare-med strict-as {zero infinity} no always-compare-med
Context	config>router>bgp>best-path-selection
Description	This command configures the comparison of BGP routes based on the MED attribute. The default behavior of SR-OS (equivalent to the no form of the command) is to only compare two routes on the basis of MED if they have the same neighbor AS (the first non-confed AS in the received AS_PATH attribute). Also by default, a route without a MED attribute is handled the same as though it had a MED attribute with the value 0. The always-compare-med command without the strict-as keyword allows MED to be compared even if the paths have a different neighbor AS; in this case, if neither zero or infinity is specified, the zero option is inferred, meaning a route without a MED is handled the same as though it had a MED attribute with the value 0. When the strict-as keyword is present, MED is only compared between paths from the same neighbor AS, and in this case, zero or infinity is mandatory and tells BGP how to interpret paths without a MED attribute.
Default	no always-compare-med

- Parameters**
- zero** — Specifies that for routes learned without a MED attribute that a zero (0) value is used in the MED comparison. The routes with the lowest metric are the most preferred.
 - infinity** — Specifies for routes learned without a MED attribute that a value of infinity ($2^{32}-1$) is used in the MED comparison. This in effect makes these routes the least desirable.
 - strict-as** — Specifies BGP paths to be compared even with different neighbor AS.

as-path-ignore

- Syntax** **as-path-ignore** [ipv4] [ipv6] [label-ipv4]
no as-path-ignore
- Context** config>service>vprn>bgp
- Description** This command configures whether AS path length is considered in the selection of the best BGP route for a prefix.
- If an address family is listed in this command, then the length of AS paths is not a factor in the route selection process for routes of that address family.
- The **no** form of the command removes the parameter from the configuration.
- Default** no as-path-ignore
- Parameters**
- ipv4** — Specifies that the AS-path length is ignored for all unlabeled unicast IPv4 routes.
 - ipv6** — Specifies that the AS-path length is ignored for all unlabeled unicast IPv6 routes.
 - label-ipv4** — Specifies that the AS-path length is ignored for all labeled-unicast IPv4 routes.

deterministic-med

- Syntax** [**no**] **deterministic-med**
- Context** config>service>vprn>bgp>best-path-selection
- Description** This command controls how the BGP decision process compares routes on the basis of MED. When **deterministic-med** is configured, BGP groups paths that are equal up to the MED comparison step based on neighbor AS, and then compares the best path from each group to arrive at the overall best path. This change to the BGP decision process makes best path selection completely deterministic in all cases. Without **deterministic-med**, the overall best path selection is sometimes dependent on the order of the route arrival because of the rule that MED cannot be compared in routes from different neighbor AS.
- Default** no deterministic-med

ebgp-ibgp-equal

Syntax	ebgp-ibgp-equal [ipv4] [ipv6] [label-ipv4] no ebgp-ibgp-equal
Context	config>router>vprn>bgp>best-path-selection
Description	This command instructs the BGP decision process to ignore the difference between EBGp and IBGP routes in selecting the best path and eligible multipaths (if multipath and ECMP are enabled). The result is a form of EIBGP load-balancing in a multipath scenario. By default (with the no form of the command), the BGP decision process prefers an EBGp learned route over an IBGP learned route. The behavior can be applied selectively to only certain types of routes by specifying one or more address family names in the command. If no families are specified, the command applies to IPv4 and IPv6 routes, and VPN-IPv4 and VPN-IPv6 routes.
Default	no ebgp-ibgp-equal
Parameters	ipv4 — Specifies that the command should be applied to unlabeled unicast IPv4 routes. ipv6 — Specifies that the command should be applied to unlabeled unicast IPv6 routes. label-ipv4 — Specifies that the command should be applied to labeled IPv4 routes.

as-override

Syntax	[no] as-override
Context	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command replaces all instances of the peer's AS number with the local AS number in a BGP route's AS_PATH. This command breaks BGP's loop detection mechanism. It should be used carefully.
Default	no as-override

authentication-key

Syntax	authentication-key [authentication-key hash-key] [hash hash2] no authentication-key
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor

Description	This command configures the BGP authentication key. Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD-5 message-based digest. The authentication key can be any combination of letters or numbers from 1 to 16. The no form of the command removes the authentication password from the configuration and effectively disables authentication.
Default	no authentication-key
Parameters	<i>authentication-key</i> — The authentication key. The key can be any combination of ASCII characters up to 255 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" "). <i>hash-key</i> — The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" "). This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided. hash — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified. hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the hash2 encrypted variable cannot be copied and pasted. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.

auth-keychain

Syntax	auth-keychain <i>name</i>
Context	<pre>config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor</pre>
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	<i>name</i> — Specifies the name of an existing keychain, up to 32 characters, to use for the specified TCP session or sessions.

best-path-selection

Syntax	best-path-selection
Context	config>service>vprn>bgp
Description	This command enables path selection configuration.

ignore-nh-metric

Syntax	ignore-nh-metric no ignore-nh-metric
Context	config>router>bgp>best-path-selection config>service>vprn config>service>vprn>bgp>best-path-selection
Description	This command instructs BGP to disregard the resolved distance to the BGP next-hop in its decision process for selecting the best route to a destination. When configured in the config>router>bgp>best-path-selection context, this command applies to the comparison of two BGP routes with the same NLRI learned from base router BGP peers. When configured in the config>service>vprn context, this command applies to the comparison of two BGP-VPN routes for the same IP prefix imported into the VPRN from the base router BGP instance. When configured in the config>service>vprn>bgp>best-path-selection context, this command applies to the comparison of two BGP routes for the same IP prefix learned from VPRN BGP peers. The no form of the command (no ignore-nh-metric) restores the default behavior whereby BGP factors distance to the next-hop into its decision process.
Default	no ignore-nh-metric

ignore-router-id

Syntax	[no] ignore-router-id
Context	config>router>bgp>best-path-selection config>service>vprn>bgp>best-path-selection
Description	When the ignore-router-id command is present and the current best path to a destination was learned from EBGp peer X with BGP identifier x and a new path is received from EBGp peer Y with BGP identifier y the best path remains unchanged if the new path is equivalent to the current best path up to the BGP identifier comparison – even if y is less than x. The no form of the command restores the default behavior of selecting the route with the lowest BGP identifier (y) as best.
Default	no ignore-router-id

bfd-enable

Syntax	[no] bfd-enable
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface. The no form of this command removes BFD from the associated BGP protocol peering.
Default	no bfd-enable

cluster

Syntax	cluster cluster-id no cluster
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command configures the cluster ID for a route reflector server. Route reflectors are used to reduce the number of IBGP sessions required within an AS. Normally, all BGP speakers within an AS must have a BGP peering with every other BGP speaker in an AS. A route reflector and its clients form a cluster. Peers that are not part of the cluster are considered to be non-clients. When a route reflector receives a route, first it must select the best path from all the paths received. If the route was received from a non-client peer, then the route reflector sends the route to all clients in the cluster. If the route came from a client peer, the route reflector sends the route to all non-client peers and to all client peers except the originator. For redundancy, a cluster can have multiple route reflectors. Confederations can also be used to remove the full IBGP mesh requirement within an AS. The no form of the command deletes the cluster ID and effectively disables the Route Reflection for the given group.
Default	no cluster — No cluster ID is defined.

Parameters *cluster-id* — The route reflector cluster ID is expressed in dot decimal notation.

Values Any 32 bit number in dot decimal notation. (0.0.0.1 to 255.255.255.255)

connect-retry

Syntax **connect-retry** *seconds*
no connect-retry

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures the BGP connect retry timer value in seconds.

When this timer expires, BGP tries to reconnect to the configured peer. This configuration parameter can be set at three levels: global level (applies to all peers), peer-group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to the default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default 120 seconds

Parameters *seconds* — Specifies the BGP connect retry timer value in seconds, expressed as a decimal integer.

Values 1 to 65535

damp-peer-oscillations

Syntax **damp-peer-oscillations** [*idle-hold-time initial-wait second-wait max-wait*] [*error-interval minutes*]

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description	This command controls how long a BGP peer session remains in the idle-state after some type of error causes the session to reset. In the idle state, BGP does not initiate or respond to attempts to establish a new session. Repeated errors that occur a short while after each session reset cause longer and longer hold times in the idle state. This command supports the DampPeerOscillations FSM behavior described in section 8.1 of RFC 4271, <i>A Border Gateway Protocol 4 (BGP-4)</i>. The default behavior, which applies when no damp-peer-oscillations is configured, is to immediately transition out of the idle-state after every reset.
Default	no damp-peer-oscillations
Parameters	<i>initial-wait</i> — Specifies the amount of time, in minutes, that a session remains in the idle-state after it has been stable for a while. Values 0 to 2048 Default 0 <i>second-wait</i> — Specifies the period of time, in minutes, that is doubled after each repeated session failure that occurs within a relatively short span of time. Values 0 to 2048 Default 5 <i>max-wait</i> — Specifies the maximum amount of time, in minutes, that a session remains in the idle-state after it has experienced repeated instability. Values 0 to 2048 Default 60 <i>minutes</i> — Specifies the interval of time, in minutes after a session reset, during which the session must be error-free in order to reset the penalty counter and return to idle-hold-time to initial-wait. Values 0 to 2048 Default 30

damping

Syntax	[no] damping
Context	<pre>config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor</pre>
Description	This command enables BGP route damping for learned routes which are defined within the route policy. Use damping to reduce the number of update messages sent between BGP peers and reduce the load on peers without affecting the route convergence time for stable routes. Damping parameters are set via route policy definition.

The **no** form of the command used at the global level disables route damping.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters:

Half-life: 15 minutes

Max-suppress: 60 minutes

Suppress-threshold: 3000

Reuse-threshold: 750

Default no damping — Learned route damping is disabled.

disable-4byte-asn

Syntax	[no] disable-4byte-asn
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command disables the use of 4-byte ASNs. It can be configured at all 3 level of the hierarchy so it can be specified down to the per peer basis. If this command is enabled 4-byte ASN support should not be negotiated with the associated remote peer(s). The no form of the command resets the behavior to the default which is to enable the use of 4-byte ASN.

disable-capability-negotiation

Syntax	[no] disable-capability-negotiation
Context	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command disables the exchange of capabilities. When command is enabled and after the peering is flapped, any new capabilities are not negotiated and strictly supports IPv4 routing exchanges with that peer.

The **no** form of the command removes this command from the configuration and restores the normal behavior.

Default no disable-capability-negotiation

disable-client-reflect

Syntax [no] disable-client-reflect

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command disables the reflection of routes by the route reflector to the group or neighbor. This only disables the reflection of routes from other client peers. Routes learned from non-client peers are still reflected to all clients.

The **no** form re-enables client reflection of routes.

Default no disable-client-reflect

disable-communities

Syntax disable-communities [standard] [extended] [large]
no disable-communities

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures BGP to disable sending standard, extended, or large communities to specific peers.

By default, all communities that are attached to a BGP route (any address family) are not stripped from the route when it is advertised to any type of peer: IBGP, EBGp or confed-EBGP.

Default no disable-communities

Parameters **standard** — Specifies that standard 4-byte communities should be removed.
extended — Specifies that 8-byte extended communities (of all types) should be removed.
large — Specifies that 12-byte large communities should be removed.

disable-fast-external-failover

Syntax	[no] disable-fast-external-failover
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command configures BGP fast external failover.

dynamic-neighbor-limit

Syntax	dynamic-neighbor-limit <i>peers</i> no dynamic-neighbor-limit
Context	config>service>vprn>bgp config>service>vprn>bgp>group
Description	This command configures the maximum number of dynamic BGP sessions that are accepted from remote peers associated with the entire BGP instance or a specific peer group. If accepting a new dynamic session would cause either the group limit or the instance limit to be exceeded, then the new session attempt is rejected and a Notification message is sent back to the remote peer. The no form of the command removes the limit on the number of dynamic sessions.
Default	no dynamic-neighbor-limit
Parameters	<i>peers</i> — Specifies the maximum number of dynamic BGP sessions.
Values	1 to 8192

dynamic-neighbor

Syntax	dynamic-neighbor
Context	config>service>vprn>bgp>group
Description	This command enters the context to configure dynamic BGP sessions for a peer group.

prefix

Syntax	[no] prefix <i>ip-prefix/prefix-length</i>
Context	config>service>vprn>bgp>group>dynamic-neighbor

Description This command configures a prefix from which to accept dynamic BGP sessions; particularly, sessions from source IP addresses not matching any configured neighbor addresses. A dynamic session is associated with the group having the longest match prefix entry for the source IP address of the peer. The group association determines local parameters that apply to the session, including the local AS, the local IP address, the peer AS, the import and export policies, and so on.

The **no** form of the command removes a prefix entry.

Parameters *ip-prefix/prefix-length* — Specifies a prefix from which to accept dynamic BGP sessions.

Values

- ipv4-prefix* — a.b.c.d (host bits must be 0)
- ipv4-prefix-length* — 0 to 32
- ipv6-prefix* — x:x:x:x:x:x:x (eight 16-bit pieces)
 x:x:x:x:x:x.d.d.d.d
 x — [0 to FFFF]H
 d — [0 to 255]D
- ipv6-prefix-length* — 0 to 128

eibgp-loadbalance

Syntax [no] eibgp-loadbalance

Context config>service>vprn>bgp

Description This command enables eiBGP load sharing so routes with both MP-BGP and IPv4 next-hops can be used simultaneously.

In order for this command to be effective, the **ecmp** and **multipath** commands for the associated VPRN instance must also be configured to allow for multiple routes to the same destination.

The **no** form of the command used at the global level reverts to default values.

Default no eibgp-loadbalance

enable-bgp-vpn-backup

Syntax enable-bgp-vpn-backup [ipv4] [ipv6]
no enable-bgp-vpn-backup

Context config>service>vprn>bgp

Description This command allows BGP-VPN routes imported into the VPRN to be used as backup paths for IPv4 and/or IPv6 BGP-learned prefixes.

Parameters **ipv4** — Allows BGP-VPN routes to be used as backup paths for IPv4 prefixes.

ipv6 — Allows BGP-VPN routes to be used as backup paths for IPv6 prefixes.

enable-peer-tracking

Syntax	[no] enable-peer-tracking
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command enables BGP peer tracking.
Default	no enable-peer-tracking

graceful-restart

Syntax	[no] graceful-restart
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command enables BGP graceful restart helper procedures (the “receiving router” role defined in the standard) for address families included in the GR capabilities of both peers. In a VPRN, SR OS can support GR helper functionality for IPv4, IPv6, label-ipv4, flow-ipv4 (IPv4 flow-spec) and flow-ipv6 (IPv6 flow-spec) routes. When a neighbor covered by the GR helper mode restarts its control plane, forwarding can continue uninterrupted while the session is re-established and routes are re-learned. The no form of the command disables graceful restart.

enable-notification

Syntax	enable-notification no enable-notification
Context	config>service>vprn>bgp>graceful-restart config>service>vprn>bgp>group>graceful-restart config>service>vprn>bgp>group>neighbor>graceful-restart
Description	When this command is present, the graceful restart capability sent by this router indicates support for NOTIFICATION messages. If the peer also supports this capability then the session can be restarted gracefully (while preserving forwarding) if either peer needs to send a NOTIFICATION message due to some type of event or error.
Default	no enable-notification

long-lived

Syntax	[no] long-lived
Context	config>service>vprn>bgp>graceful-restart config>service>vprn>bgp>group>graceful-restart config>service>vprn>bgp>group>neighbor>graceful-restart
Description	This command enables the context to configure BGP Long-Lived Graceful-Restart (LLGR) procedures. LLGR, known informally as BGP persistence, is an extension of BGP graceful restart that allows a session to stay down for a longer period of time. During this time, learned routes are marked and re-advertised as stale but they can continue to be used as routes of last resort. The LLGR handling of a session failure can be invoked immediately or it can be delayed until the end of the traditional GR restart window.
Default	no long-lived

advertise-stale-to-all-neighbors

Syntax	advertise-stale-to-all-neighbors [without-no-export no without-no-export] no advertise-stale-to-all-neighbors
Context	config>service>vprn>bgp>graceful-restart>long-lived config>service>vprn>bgp>group>graceful-restart>long-lived config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived
Description	This command allows BGP routes marked as LLGR stale to be advertised to BGP peers that did not advertise the LLGR capability when the session was opened. The no version of this command causes advertisement behavior to follow the rule that stale routes cannot be advertised to a peer that does not understand or implement the LLGR capability. Stale routes are withdrawn towards such peers. When this command is configured with the without-no-export option, LLGR stales routes can be advertised to any peer (EBGP or IBGP) that did not signal the LLGR capability. Towards IBGP and confederation-EBGP peers that did not advertise the LLGR capability, the LOCAL_PREFERENCE attribute in the advertised stale routes is automatically set to zero. When this command is configured without the without-no-export option, LLGR stale routes are not advertised to any EBGP peer that did not signal the LLGR capability. Towards IBGP and confederation-EBGP peers that did not advertise the LLGR capability the LOCAL_PREFERENCE attribute in the advertised stale routes is automatically set to zero and a NO_EXPORT standard community is automatically added to the routes.
Default	no advertise-stale-to-all-neighbors
Parameters	without-no-export — Allows LLGR stale routes to be advertised to all peers, such that they can exit the local AS.

advertised-stale-time

Syntax	advertised-stale-time <i>seconds</i> no advertised-stale-time
Context	config>service>vprn>bgp>graceful-restart>long-lived config>service>vprn>bgp>graceful-restart>long-lived>family config>service>vprn>bgp>group>graceful-restart>long-lived config>service>vprn>bgp>group>graceful-restart>long-lived>family config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived>family
Description	This command sets the value of the long-lived stale time that is advertised by the router in its LLGR capability. When configured in the long-lived configuration context, advertised-stale-time applies to all AFI/SAFI in the advertised LLGR capability except for any AFI/SAFI with a family-specific override. A family-specific override is configured with the advertised-stale-time command in a family context. The no version of this command sets the advertised-stale-time value to 24 hours (86400 seconds).
Default	no advertised-stale-time
Parameters	<i>seconds</i> — Specifies the advertised long-lived stale time in seconds. Values 0 to 16777215

family

Syntax	[no] family { ipv4 ipv6 label-ipv4 flow-ipv4 flow-ipv6 }
Context	config>service>vprn>bgp>graceful-restart>long-lived config>service>vprn>bgp>group>graceful-restart>long-lived config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived
Description	This command configures family-specific LLGR parameters for BGP peers.
Default	no family
Parameters	ipv4 — Specifies the IPv4 family. ipv6 — Specifies the IPv6 family. label-ipv4 — Specifies the label IPv4 family. flow-ipv4 — Specifies the flow IPv4 family. flow-ipv6 — Specifies the flow IPv6 family.

helper-override-stale-time

Syntax	helper-override-stale-time <i>seconds</i> no helper-override-stale-time
Context	config>service>vprn>bgp>graceful-restart>long-lived config>service>vprn>bgp>graceful-restart>long-lived>family config>service>vprn>bgp>group>graceful-restart>long-lived config>service>vprn>bgp>group>graceful-restart>long-lived>family config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived>family
Description	This command overrides the LLGR stale-time advertised by a peer (in its LLGR capability) with a locally-configured value. When configured in the long-lived configuration context, helper-override-stale-time applies to all AFI/SAFI in the advertised LLGR capability except for any AFI/SAFI with a family-specific override. A family-specific override is configured with the helper-override-stale-time command in a family context. By default, the LLGR stale-time for an AFI/SAFI is the value signaled by the peer in the corresponding AFI/SAFI part of the LLGR capability.
Default	no helper-override-stale-time
Parameters	<i>seconds</i> — Specifies the locally imposed LLGR stale time in seconds.
	Values 0 to 16777215

forwarding-bits-set

Syntax	forwarding-bits-set {all non-fwd} no forwarding-bits-set
Context	config>service>vprn>bgp>graceful-restart>long-lived config>service>vprn>bgp>group>graceful-restart>long-lived config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived
Description	This command determines the setting of the F bits in the GR and LLGR capabilities advertised by the router. When the F bit is set for an AFI/SAFI, it indicates that the advertising router was able to preserve forwarding state for the routes of that AFI/SAFI across the last restart. If a router restarts and does not set F=1, then when the session with a peer is re-established, the peer immediately deletes all LLGR stale routes it was preserving on behalf of the restarting router for the corresponding AFI/SAFI. This command allows the F bits for all advertised AFI/SAFI to be set to 1, or only the F bits for non-forwarding AFI/SAFI to be set to 1. Non-forwarding AFI/SAFI are the following configuration-related address families: L2-VPN, route-target, flow-IPv4, and flow-IPv6.
Default	no forwarding-bits-set
Parameters	all — Specifies that the F bit for all AFI/SAFI should be set to 1.

non-fwd — Specifies that the F bit for only non-forwarding AFI/SAFI should be set to 1. These AFI/SAFI correspond to the following families: L2-VPN, route-target, flow-IPv4, and flow-IPv6.

helper-override-restart-time

Syntax	helper-override-restart-time <i>seconds</i> no helper-override-restart-time
Context	config>service>vprn>bgp>graceful-restart>long-lived config>service>vprn>bgp>group>graceful-restart>long-lived config>service>vprn>bgp>group>neighbor>graceful-restart>long-lived
Description	This command overrides the restart-time advertised by a peer (in its GR capability) with a locally-configured value. This override applies only to AFI/SAFI that were included in the GR capability of the peer. The restart-time is always zero for AFI/SAFI not included in the GR capability. This command is useful if the local router wants to force LLGR phase to begin after a set time for all protected AFI/SAFI. By default, the restart time for all AFI/SAFI in the GR capability is the value signaled by the peer.
Default	no helper-override-restart-time
Parameters	<i>seconds</i> — The locally-imposed restart time for all AFI/SAFI included in the peer's GR capability. Values 0 to 4095

restart-time

Syntax	restart-time <i>seconds</i> no restart-time
Context	config>service>vprn>bgp>graceful-restart config>service>vprn>bgp>group>graceful-restart config>service>vprn>bgp>group>neighbor>graceful-restart
Description	This command sets the value of the restart-time that is advertised in the router's graceful-restart capability. If this command is not configured, the default is 300.
Default	no restart-time
Parameters	<i>seconds</i> — Specifies the restart-time that is advertised in the router's graceful-restart capability. Values 0 to 4095 seconds Default 300

stale-routes-time

Syntax	[no] stale-routes-time <i>time</i>
Context	config>service>vprn>bgp>graceful-restart config>service>vprn>bgp>group>graceful-restart config>service>vprn>bgp>group>neighbor>graceful-restart
Description	This command configures the time period to keep stale routes before the END-OF-RIB message is received from the restarting router.
Default	360 seconds
Parameters	<i>time</i> — [1 to 3600 seconds]

enforce-first-as

Syntax	enforce-first-as
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	When this command is configured so that it applies to an EBGp session, all routes (belonging to all address families) that are received from the EBGp peer are checked to ensure that the most recent autonomous system number (ASN) in the AS_PATH attribute of each route matches the configured peer-as of the session; if it does not match, then either the session is reset (if update-fault-tolerance is not enabled) or the session is left up but the route is treated as withdrawn (if update-fault-tolerance is enabled). Enabling or disabling this command on a session that is already up does not flap the session. When enforce-first-as is enabled, previously received routes are not checked for compliance with the rule. Enforcement applies only to routes received after the command is enabled and stops when the command is disabled.

error-handling

Syntax	error-handling
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command specifies whether the error handling mechanism for optional transitive path attributes is enabled for this peer group.

update-fault-tolerance

Syntax	[no] update-fault-tolerance
Context	config>service>vprn>bgp>error-handling config>service>vprn>bgp>group>error-handling config>service>vprn>bgp>group>neighbor>error-handling
Description	This command enables treat-as-withdraw and other similarly non-disruptive approaches for handling a wide range of UPDATE message errors, as long as there are no length errors that prevent all of the NLRI fields from being correctly identified and parsed.
Default	no update-fault-tolerance

export

Syntax	export <i>plcy-or-long-expr</i> [<i>plcy-or-expr</i> [<i>plcy-or-expr</i> ... (up to 14 max)]] no export
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command is used to specify route policies that control how outbound routes transmitted to certain peers are handled. Route policies are configured in the config>router>policy-options context. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in a peer-group) or neighbor level (only applies to the specified peer). The most specific level is used. The export command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine the modifications of each route and the final action to accept or reject the route. Only one of the 15 objects referenced by the export command can be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters while the remaining 14 objects have a maximum length of 64 characters each. When multiple export commands are issued, the last command entered overrides the previous command. When an export policy is not specified, BGP-learned routes are advertised by default; non-BGP routes are not advertised. The no form of this command removes the policy association.
Default	no export

- Parameters** *plcy-or-long-expr* — Specifies the route policy name, up to 64 characters in length, or a policy logical expression, up to 255 characters in length.
- plcy-or-expr* — Specifies the route policy name, up to 64 characters in length, or a policy logical expression, up to 255 characters in length.

family

- Syntax** **family [ipv4] [label-ipv4] [ipv6] [mcast-ipv4] [flow-ipv4] [mcast-ipv6] [flow-ipv6]**
no family
- Context** config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
- Description** This command configures the set of BGP address families (AFI plus SAFI) to be supported by the applicable VPRN BGP sessions.
- The **no** form of the command restores the default, which corresponds to unlabeled IPv4 unicast routes (AFI 1, SAFI 1) only.
- Default** family ipv4
- Parameters** **ipv4** — Adds support for the IPv4 unicast (unlabeled) address family.
label-ipv4 — Adds support for the IPv4 unicast (labeled) address family.
ipv6 — Adds support for the IPv6 unicast (unlabeled) address family.
mcast-ipv4 — Adds support for the IPv4 multicast SAFI address family.
flow-ipv4 — Adds support for the IPv4 flow-spec address family.
mcast-ipv6 — Adds support for the IPv4 multicast SAFI address family.
flow-ipv6 — Adds support for the IPv4 flow-spec address family.

flowspec

- Syntax** **flowspec**
- Context** config>service>vprn>bgp
- Description** The context to enable and disable flowspec validations.

validate-dest-prefix

- Syntax** **validate-dest-prefix**
no validate-dest-prefix

Context	config>service>vprn>bgp>flowspec
Description	This command enables or disables validation of received IPv4 and IPv6 flowspec routes that contain a destination-prefix subcomponent. A flowspec route with a destination-prefix subcomponent is considered invalid if both of the following are true: • it was originated outside the local AS of the receiving BGP router • the neighbor AS of the flowspec route does not match the neighbor AS of the best match BGP (unicast) route for the destination prefix or the neighbor AS of any longer match BGP (unicast) route for the destination prefix An invalid route is retained in the BGP but it is not used for filtering traffic or propagated to other BGP routers. The no form of the command disables the validation procedure based on destination-prefix.
Default	no validate-dest-prefix

validate-redirect-ip

Syntax	validate-redirect-ip no validate-redirect-ip
Context	config>service>vprn>bgp>flowspec
Description	This command enables procedures to validate the redirect-to-IPv4 action attached to flowspec-IPv4 routes received by the BGP instance. The SR OS flow-spec implementation supports the redirect-to-IPv4 action encoded as an IPv4-address-specific BGP extended community. When this command is configured, a flowspec-IPv4 route is considered invalid and not installed as a filter rule if the flowspec-IPv4 route is deemed to have originated in a different AS than the IP route that resolves the redirection IPv4 address. The originating AS of a flow-spec route is determined from its AS path. The no form of the command disables the check described above.
Default	no validate-redirect-ip

group

Syntax	group <i>name</i> [<i>esm-dynamic-peer</i>] no group
Context	config>service>vprn>bgp

Description	This command creates a context to configure a BGP peer group. The no form of the command deletes the specified peer group and all configurations associated with the peer group. The group must be shut down before it can be deleted.
Parameters	<i>name</i> — The peer group name. Allowed values is a string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. esm-dynamic-peer — This flag designates that the given BGP group is used by BGP peers created dynamically based on subscriber-hosts pointing to corresponding BGP peering policy. There can be only one BGP group with this flag set in any given VPRN. No BGP neighbors can be manually configured in a BGP group with this flag set. Default disabled

neighbor

Syntax	[no] neighbor <i>ip-address</i>												
Context	config>service>vprn>bgp>group												
Description	This command creates a BGP peer/neighbor instance within the context of the BGP group. This command can be issued repeatedly to create multiple peers and their associated configuration. The no form of the command is used to remove the specified neighbor and the entire configuration associated with the neighbor. The neighbor must be administratively shutdown before attempting to delete it. If the neighbor is not shut down, the command will not result in any action except a warning message on the console indicating that neighbor is still administratively up.												
Parameters	<i>ip-address</i> — The IP address of the BGP peer router in dotted decimal notation. Values <table> <tr> <td>ipv4-address</td><td>a.b.c.d (host bits must be 0)</td></tr> <tr> <td>ipv6-address</td><td>x:x:x:x:x:x:x[-interface]</td></tr> <tr> <td></td><td>x:x:x:x:x:x.d.d.d.d[-interface]</td></tr> <tr> <td></td><td>x: [0 to FFFF]H</td></tr> <tr> <td></td><td>d: [0 to 255]D</td></tr> <tr> <td></td><td>interface: 32 characters maximum, mandatory for link local addresses</td></tr> </table>	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: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
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: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												

The ipv6-address applies to the 7750 SR only.

family

Syntax	family [ipv4] [ipv6] [mcast-ipv4] no family
Context	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command specifies the address family or families to be supported over BGP peerings in the base router. This command is additive so issuing the family command adds the specified address family to the list. The no form of the command removes the specified address family from the associated BGP peerings. If an address family is not specified, then reset the supported address family back to the default.
Default	ipv4
Parameters	ipv4 — Provisions support for IPv4 routing information. ipv6 — Exchange IPv6 routing information (applies to the 7750 SR only). mcast-ipv4 — Provisions Multicast IPv4 support.

hold-time

Syntax	hold-time <i>seconds</i> [<i>min seconds2</i>] no hold-time
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command configures the BGP hold time, expressed in seconds. The BGP hold time specifies the maximum time BGP waits between successive messages (either keepalive or update) from its peer, before closing the connection. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used. Even though the router OS implementation allows setting the keepalive time separately, the configured keepalive timer is overridden by the hold-time value under the following circumstances: 1. If the specified hold-time is less than the configured keepalive time, then the operational keepalive time is set to a third of the hold-time; the configured keepalive time is not changed.

2. If the **hold-time** is set to zero, then the operational value of the **keepalive** time is set to zero; the configured **keepalive** time is not changed. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of the command used at the global level reverts to the default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default	90 seconds
Parameters	<i>seconds</i> — 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
	<i>seconds2</i> — Specifies the minimum hold-time that is accepted for the session. If the peer proposes a hold-time lower than this value the session attempt is rejected.

ibgp-multipath

Syntax	[no] ibgp-multipath
Context	config>service>vprn>bgp
Description	This command defines the type of IBGP multipath to use when adding BGP routes to the route table if the route resolving the BGP nexthop offers multiple next-hops. The no form of the command disables the IBGP multipath load balancing feature.

import

Syntax	import <i>plcy-or-long-expr</i> [<i>plcy-or-expr</i> [<i>plcy-or-expr</i> ... (up to 14 max)]] no import
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command is used to specify route policies that control the handling of inbound routes received from certain peers. Route policies are configured in the config>router>policy-options context. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in a peer-group) or neighbor level (only applies to the specified peer). The most specific level is used

The **import** command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine the modifications of each route and the final action to accept or reject the route.

Only one of the 15 objects referenced by the **import** command can be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters while the remaining 14 objects have a maximum length of 64 characters each.

When multiple **import** commands are issued, the last command entered overrides the previous command.

When an import policy is not specified, BGP routes are accepted by default.

The **no** form of this command removes the policy association.

Default	no import
Parameters	<i>plcy-or-long-expr</i> — The route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long). <i>plcy-or-expr</i> — The route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long).

initial-send-delay-zero

Syntax	[no] initial-send-delay-zero
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command configures BGP to send UPDATE messages announcing reachability information to a peer or set of peers immediately after the sessions come up (become established) with these peers. The default behavior, provided by the no form of this command, is to wait for min-route-advertisement time after each session is established before sending the first set of UPDATE messages.

keepalive

Syntax	keepalive seconds no keepalive
Context	config>service>vprn>bgp config>service>vprn>bgp>group

	config>service>vprn>bgp>group>neighbor
Description	This command configures the BGP keepalive timer. A keepalive message is sent every time this timer expires. The seconds parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used. The keepalive value is generally one-third of the hold-time interval. Even though the OS implementation allows the keepalive value and the hold-time interval to be independently set, under the following circumstances, the configured keepalive value is overridden by the hold-time value: If the specified keepalive value is greater than the configured hold-time, then the specified value is ignored, and the keepalive is set to one third of the current hold-time value. If the specified hold-time interval is less than the configured hold-time value, then the keepalive value is reset to one third of the specified hold-time interval. If the hold-time interval is set to zero, then the configured value of the keepalive value is ignored. This means that the connection with the peer is up permanently and no keepalive packets are sent to the peer. The no form of the command used at the global level reverts to the default value. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	30 seconds
Parameters	seconds — The keepalive timer in seconds, expressed as a decimal integer. Values 0 to 21845

label-preference

Syntax	label-preference <i>value</i> no label-preference
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command configures the route preference for routes learned from labeled-unicast peers. This command can be configured at three levels: • Global level — applies to all peers

- Group level — applies to all peers in the peer-group
- Neighbor level — applies only to the specified peer

The most specific value is used.

The lower the preference, the higher the chance of the route being the active route.

The **no** form of the command used at the global level reverts to the default *value* of 170.

The **no** form of the command used at the group level reverts to the *value* defined at the global level.

The **no** form of the command used at the neighbor level reverts to the *value* defined at the group level.

Default	no label-preference
Parameters	<i>value</i> — Specifies the route preference value.
Values	1 to 255

link-bandwidth

Syntax	link-bandwidth
Context	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command enables the configuration context for handling the link-bandwidth extended community attached to specific BGP routes. When all used multipaths of an IP prefix correspond to BGP routes with a link-bandwidth extended community, the datapath is programmed to do weighted ECMP across the BGP next-hops in proportion to the bandwidth values.

accept-from-ebgp

Syntax	accept-from-ebgp <i>family</i> [<i>family</i>] no accept-from-ebgp
Context	config>service>vprn>bgp>group>link-bandwidth config>service>vprn>bgp>group>neighbor>link-bandwidth
Description	This command configures BGP to accept and use the link-bandwidth extended community attached to any route received from any EBGp peer in the scope of the command, as long as that route belongs to one of the listed address families.

The link-bandwidth extended community is encoded as a non-transitive type. This means that by default it should not be attached to any route advertised to an EBGp peer and it should be discarded when received in any route from an EBGp peer. This command overrides the standard behavior.

Up to three families may be configured.

The **no** form of this command restores the default behavior of discarding the link-bandwidth extended community in any route received from an EBGp peer.

Default	no accept-from-ebgp
Parameters	<i>family</i> — Specifies the address families for which receiving the link-bandwidth extended community from EBGp peers should be supported.
Values	ipv4 — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes. label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes. ipv6 — Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes.

add-to-received-ebgp

Syntax	add-to-received-ebgp <i>family</i> [<i>family</i>] no add-to-received-ebgp
Context	config>service>vprn>bgp>group>link-bandwidth config>service>vprn>bgp>group>neighbor>link-bandwidth
Description	This command configures BGP to automatically add a link-bandwidth extended community to every route received from a directly connected (single-hop) EBGp peer within the scope of the command, as long as that route belongs to one of the listed address families. The link-bandwidth extended community added by this command encodes the local-AS number of receiving BGP instance and the bandwidth of the interface to the directly connected EBGp peer. Up to three families may be configured. The no form of this command removes the link-bandwidth extended community added to received BGP routes.
Default	no add-to-received-ebgp
Parameters	<i>family</i> — Specifies the address families for which receiving the link-bandwidth extended community from EBGp peers should be supported.
Values	ipv4 — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes.

label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes.

ipv6 — Adds a link-bandwidth extended community to unlabeled-unicast IPv6 routes.

aggregate-used-paths

Syntax	aggregate-used-paths <i>family</i> [<i>family</i>] no aggregate-used-paths
Context	config>service>vprn>bgp>group>link-bandwidth config>service>vprn>bgp>group>neighbor>link-bandwidth
Description	This command configures BGP to aggregate the bandwidth values from the link-bandwidth extended communities of the used multipaths towards an IP prefix when it is re-advertising a route with next-hop-self towards peers within the scope of the command, as long as the route belongs to one of the listed address families. Aggregation is not supported unless all of the used multipaths (up to the configured ECMP limit) correspond to received BGP routes with a link-bandwidth extended community. If add-path is also enabled toward the peer, then all of the add-paths advertised to the peer encode the aggregated bandwidth in a link-bandwidth extended community. Up to three families may be configured. The no form of the command disables aggregation in a next-hop-self scenario and the link-bandwidth extended community in the advertised route is a copy of the link-bandwidth extended community in the received route (which may have been added by import policy or by the effect of the add-to-received-ebgp command).
Default	no aggregate-used-paths
Parameters	<i>family</i> — Specifies the address families for which receiving the link-bandwidth extended community from EBGp peers should be supported. Values ipv4 — Adds a link-bandwidth extended community to unlabeled-unicast IPv4 routes. label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes. ipv6 — Adds a link-bandwidth extended community to unlabeled-unicast IPv6 routes.

send-to-ebgp

Syntax	send-to-ebgp <i>family</i> [<i>family</i>] no send-to-ebgp

Context	config>service>vprn>bgp>group>link-bandwidth config>service>vprn>bgp>group>neighbor>link-bandwidth
Description	This command configures BGP to allow link-bandwidth extended community to be sent in routes advertised to EBGp peers in the scope of the command, as long the routes belong to one of the listed address families. The link-bandwidth extended community is encoded as a non-transitive type. This means that by default it should not be attached to any route advertised to an EBGp peer and it should be discarded when received in any route from an EBGp peer. This command overrides the standard behavior. Up to three families may be configured. The no form of the command restores the default behavior of stripping the link-bandwidth extended community from any route advertised to an EBGp peer.
Default	no send-to-ebgp
Parameters	<i>family</i> — Specifies the address families for which receiving the link-bandwidth extended community from EBGp peers should be supported. Values - ipv4 — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes. - label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes. - ipv6 — Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes.

local-address

Syntax	local-address <i>ip-address</i> no local-address
Context	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	Configures the local IP address used by the group or neighbor when communicating with BGP peers. Outgoing connections use the local-address as the source of the TCP connection when initiating connections with a peer. When a local address is not specified, the OS uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGp peers. This command is used at the neighbor level to revert to the value defined under the group level. The no form of the command removes the configured local-address for BGP.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Parameters **no local-address** — The router ID is used when communicating with IBGP peers and the interface address is used for directly connected EBGP peers.

ip-address — The local address expressed in dotted decimal notation. Allowed values are a valid routable IP address on the router, either an interface or system IP address.

local-as

Syntax **local-as** *as-number* [**private**] [**no-prepend-global-as**]
no local-as

Context config>service>vprn>bgp
 config>service>vprn>bgp>group
 config>service>vprn>bgp>group>neighbor

Description This command configures a BGP virtual autonomous system (AS) number.

In addition to the global AS number configured for BGP in the config>router>autonomous-system context, a virtual (local) AS number can be configured to support various AS number migration scenarios. The local AS number is added to the beginning of the as-path attribute ahead of the router's AS number.

This configuration parameter can be set at three levels: global level (applies to all EBGP peers), group level (applies to all EBGP peers in peer-group) or neighbor level (only applies to EBGP specified peer). Thus, by specifying this at each neighbor level, it is possible to have a separate local-as per EBGP session. The local-as command is not supported for IBGP sessions. When the optional **private** keyword is specified in the command the local-as number is not added to inbound routes from the EBGP peer that has **local-as** in effect.

When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The **private** attribute can be added or removed dynamically by reissuing the command.

Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local AS number. Changing the local AS at the global level in an active BGP instance causes BGP to re-establish the peer relationships with all peers in the group with the new local AS number. Changing the local AS at the neighbor level in an active BGP instance causes BGP to re-establish the peer relationship with the new local AS number.

This is an optional command and can be used in the following circumstance:

Provider router P is moved from AS1 to AS2. The customer router that is connected to P, however, is configured to belong to AS1. To avoid reconfiguring the customer router, the **local-as** value on router P can be set to AS1. Thus, router P adds AS1 to the as-path message for routes it advertises to the customer router.

The **no** form of the command used at the global level removes any virtual AS number configured.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default	no local-as
Parameters	<i>as-number</i> — The virtual autonomous system number, expressed as a decimal integer.
Values	1 to 65535
	private — Specifies the local-as is hidden in paths learned from the peering.
	no-prepend-global-as — Specifies that the global-as is hidden in paths announced to the EBGp peer.

local-preference

Syntax	local-preference <i>local-preference</i> no local-preference
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
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 configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used. The no form of the command at the global level specifies that incoming routes with local-preference set are not overridden and routes arriving without local-preference set are interpreted as if the route had local-preference value of 100. The no form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default	no local-preference - Does not override the local-preference value set in arriving routes and analyze routes without local preference with value of 100.
Parameters	<i>local-preference</i> — The local preference value to be used as the override value, expressed as a decimal integer.
Values	0 to 4294967295

loop-detect

Syntax	loop-detect {drop-peer discard-route ignore-loop off} no loop-detect
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command configures how the BGP peer session handles loop detection in the AS path. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in 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 the command used at the global level reverts to default, which is loop-detect ignore-loop. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	loop-detect ignore-loop
Parameters	drop-peer — Sends a notification to the remote peer and drops the session. discard-route — Discards routes received with loops in the AS path. ignore-loop — ignores routes with loops in the AS path but maintains peering. off — Disables loop detection.

med-out

Syntax	med-out {number igp-cost}

	no med-out
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set. The specified value can be overridden by any value set via a route policy. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used. The no form of the command used at the global level reverts to default where the MED is not advertised. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	no med-out
Parameters	<i>number</i> — The MED path attribute value, expressed as a decimal integer. Values 0 to 4294967295 igp-cost — The MED is set to the IGP cost of the given IP prefix.

min-route-advertisement

Syntax	min-route-advertisement <i>seconds</i> no min-route-advertisement
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command configures the minimum interval, in seconds, at which a prefix can be advertised to a peer. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used. The no form of the command reverts to default values.
Default	30 seconds

Parameters	<i>seconds</i> — The minimum route advertising interval, in seconds, expressed as a decimal integer.
Values	1 to 255

multihop

Syntax	multihop <i>ttl-value</i> no multihop
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command configures the time to live (TTL) value entered in the IP header of packets sent to an EBGp peer multiple hops away. This parameter is meaningful only when configuring EBGp peers. It is ignored if set for an IBGP peer. The no form of the command is used to convey to the BGP instance that the EBGp peers are directly connected. The no form of the command reverts to default values.
Default	1 (EBGP peers are directly connected) 64 (IBGP)
Parameters	<i>ttl-value</i> — The TTL value, expressed as a decimal integer.
Values	1 to 255

multipath

Syntax	multipath <i>max-paths</i> [ebgp <i>ebgp-max-paths</i>] [ibgp <i>ibgp-max-paths</i>] [restrict { same-neighbor-as exact-as-path }] no multipath
Context	config>service>vprn>bgp
Description	This command enables BGP multipath for the IPv4 and IPv6 address families. When multipath is enabled, traffic to the destination is load-shared across a set of paths (BGP routes) that the BGP decision process considers 'equal' to the best path. The actual distribution of traffic over the multiple paths may be equal or unequal (that is, based on weights derived from the Link Bandwidth extended community).

To qualify as a multipath, a non-best route must meet the following criteria (some criteria are controlled by this command):

- must be the same type of route as the best path (same AFI/SAFI and, in some cases, same next-hop resolution method)
- must be tied with the best path for all criteria of equal or higher importance than IGP cost to reach the BGP next-hop, except for criteria that are configured to be ignored
- must not have the same BGP next-hop as the best path or any other multipath
- must not cause the ECMP limit of the routing instance to be exceeded (configurable using the **ecmp** command to a value in the range 1-64)
- must not cause the configured *max-paths* limit of the BGP instance to be exceeded
If the best path is an EBGp learned route and the **ebgp** option is present, the *ebgp-max-paths* limit overrides the *max-paths* limit. If the best path is an IBGP learned route and the **ibgp** option is present, the *ibgp-max-paths* limit overrides the *max-paths* limit.
All path limits are configurable up to a maximum of 64. Multipath is effectively disabled if a value is set to one
- must have the same neighbor AS in its AS path as the best path if the **restrict same-neighbor-as** option is configured
By default, any path with the same AS path length as the best path (regardless of neighbor AS) is eligible for multipath.
- must have the exact same AS path as the best path if the **restrict exact-as-path** option is configured
By default, any path with the same AS path length as the best path (regardless of the actual AS numbers) is eligible for multipath.

The **no** form of the command disables BGP multipath (equivalent to multipath 1).

Default	no multipath
Parameters	<i>max-paths</i> — The maximum number of multipaths per prefix/NLRI if <i>ebgp-max-paths</i> or <i>ibgp-max-paths</i> does not apply. Values 1 to 64 <i>ebgp-max-paths</i> — The maximum number of multipaths per prefix/NLRI when the best path is an EBGp-learned route. Values 1 to 64 <i>ibgp-max-paths</i> — The maximum number of multipaths per prefix/NLRI when the best path is an IBGP-learned route. Values 1 to 64 same-neighbor-as — The non-best path must have the same neighbor AS in its AS path as the best path. exact-as-path-as — The non-best path must have the exact same AS path as the best path.

next-hop-resolution

Syntax	next-hop-resolution
Context	config>service>vprn>bgp
Description	This command enters the context to configure next-hop resolution parameters.

next-hop-self

Syntax	[no] next-hop-self
Context	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command configures the group or neighbor to always set the NEXTHop path attribute to its own physical interface when advertising to a peer. This is primarily used to avoid third-party route advertisements when connected to a multi-access network. The no form of the command used at the group level allows third-party route advertisements in a multi-access network. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Parameters	no next-hop-self — Third-party route advertisements are allowed.

passive

Syntax	[no] passive
Context	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command enables passive mode for the BGP group or neighbor. When in passive mode, BGP will not attempt to actively connect to the configured BGP peers but responds only when it receives a connect open request from the peer. The no form of the command used at the group level disables passive mode where BGP actively attempts to connect to its peers. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Parameters	no passive — BGP will actively try to connect to all the configured peers.

peer-as

Syntax	peer-as <i>as-number</i>
Context	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer. For EBGp peers, the peer AS number configured must be different from the autonomous system number configured for this router under the global level since the peer will be in a different autonomous system than this router For IBGP peers, the peer AS number must be the same as the autonomous system number of this router configured under the global level. This is a required command for each configured peer. This may be configured under the group level for all neighbors in a particular group.
Default	No AS numbers are defined.
Parameters	<i>as-number</i> — The autonomous system number, expressed as a decimal integer.
Values	1 to 65535

policy

Syntax	policy <i>policy-name</i> no policy
Context	config>service>vprn>bgp>next-hop-res
Description	This command specifies the name of a policy statement to use with the BGP next-hop resolution process. The policy controls which IP routes in RTM are eligible to resolve the BGP next-hop addresses of IPv4 and IPv6 routes. The policy has no effect on the resolution of BGP next-hops to MPLS tunnels. If a BGP next-hop of an IPv4 or IPv6 route R is resolved in RTM and the longest matching route for the next-hop address is an IP route N that is rejected by the policy then route R is unresolved; if the route N is accepted by the policy then it becomes the resolving route for R. The default next-hop resolution policy (when the no policy command is configured) is to use the longest matching active route in RTM that is not a BGP route (unless use-bgp-routes is configured), an aggregate route or a subscriber management route.
Default	no policy

Parameters *policy-name* — The route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. Route policies are configured in the **config>router>policy-options** context.

peer-tracking-policy

Syntax **peer-tracking-policy** *policy-name*
no peer-tracking-policy

Context config>service>vprn>bgp

Description This command specifies the name of a policy statement to use with the BGP peer-tracking function on the BGP sessions where this is enabled. The policy controls which IP routes in RTM are eligible to indicate reachability of IPv4 and IPv6 BGP neighbor addresses. If the longest matching route in RTM for a BGP neighbor address is an IP route that is rejected by the policy, or it is a BGP route accepted by the policy, or if there is no matching route, the neighbor is considered unreachable and BGP tears down the peering session and holds it in the idle state until a valid route is once again available and accepted by the policy.

The default peer-tracking policy (when the **no peer-tracking-policy** command is configured) is to use the longest matching active route in RTM that is not an LDP shortcut route or an aggregate route.



Note: When **peer-tracking** is configured, the peer-tracking policy should only permit one of **direct-interface** or **direct** routes to be advertised to a BGP peer. Advertising both routes will cause the best route to oscillate.

Default no peer-tracking-policy

Parameters *policy-name* — The route policy name. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. Route policies are configured in the **config>router>policy-options** context.

preference

Syntax [**no**] **preference** *preference*

Context config>service>vprn>bgp
config>service>vprn>bgp>group

Description This command configures the route preference for routes learned from the configured peer(s).

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The lower the preference the higher the chance of the route being the active route. The OS assigns BGP routes highest default preference compared to routes that are direct, static or learned via MPLS or OSPF.

The **no** form of the command used at the global level reverts to default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default	170
Parameters	<i>preference</i> — The route preference, expressed as a decimal integer.
Values	1 to 255

prefix-limit

Syntax	prefix-limit <i>family limit</i> [log-only] [threshold percentage] [idle-timeout { <i>minutes</i> forever }] [log-only] [post-import] no prefix-limit <i>family</i>
Context	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command configures the maximum number of BGP routes that can be received from a peer before administrative action is taken. The administrative action can be the generation of a log event or taking down the session. If a session is taken down, then it can be brought back up automatically after an idle-timeout period, or else it can be configured to stay down ('forever') until the operator performs a reset. The prefix-limit command allows each address family to have its own limit; a set of address family limits can be applied to one neighbor or to all neighbors in a group. The no form of the command removes the prefix-limit .
Default	No prefix limits for any address family.
Parameters	<i>percent</i> — The threshold value (as a percentage) that triggers a warning message to be sent. Values 1 to 100 <i>family</i> — The address family to which the limit applies. Values ipv4, label-ipv4, ipv6, mcast-ipv4, flow-ipv4, flow-ipv6, mcast-ipv6

limit — The number of routes that can be learned from a peer expressed as a decimal integer.

Values 1 to 4294967295

minutes — Specifies duration in minutes before automatically re-establishing a session.

Values 1 to 1024

forever — Specifies that the session is reestablished only after **clear router bgp** command is executed.

log-only — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is reached. However, the BGP session is not taken down.

post-import — Specifies that the limit should be applied only to the number of routes that are accepted by import policies.

rapid-withdrawal

Syntax [no] **rapid-withdrawal**

Context config>service>vprn>bgp

Description This command disables the delay (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 the command removes this command from the configuration and returns withdrawal processing to the normal behavior.

Default no rapid-withdrawal

remove-private

Syntax **remove-private** [limited] [skip-peer-as]
no remove-private

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command allows private AS numbers to be removed from the AS path before advertising them to BGP peers.

When the **remove-private** parameter is set at the global level, it applies to all peers regardless of group or neighbor configuration. When the parameter is set at the group level, it applies to all peers in the group regardless of the neighbor configuration.

The set of AS numbers that are defined by IANA as private are in the range of 64512 through 65535, and 4200000000-4294967295, inclusive.

The **no** form of the command used at the global level reverts to default value. The **no** form of the command used at the group level reverts to the value defined at the global level. The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

- Default** no remove-private - Private AS numbers will be included in the AS path attribute.
- Parameters** *limited* — This optional keyword removes private ASNs up to the first public ASN encountered. It then stops removing private ASNs.
- skip-peer-as* — This optional keyword causes this command to not remove a private ASN from the AS-Path if that ASN is the same as the BGP peer AS number.

rib-management

- Syntax** **rib-management**
- Context** config>service>vprn>bgp
- Description** This command enters the context to configure RIB management parameters.

leak-import

- Syntax** **leak-import** *plcy-or-long-expr* [*plcy-or-expr* ... (up to 14 max)]
no leak-import
- Context** config>service>vprn>bgp>rib-management>ipv4
config>service>vprn>bgp>rib-management>label-ipv4
config>service>vprn>bgp>rib-management>ipv6
- Description** This command is used to specify route policies that control the importation of leak-eligible routes from the BGP RIB of another routing instance into the unlabeled-IPv4, unlabeled-IPv6, or labeled-IPv4 RIB of the base router. To leak a route from one routing instance to another, the origin and destination RIB types must be the same; for example, it is not possible to leak a route from an unlabeled-IPv4 RIB of a VPRN into the labeled-IPv4 RIB of the base router.
- The **leak-import** command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine final action to accept or reject the route.
- Only one of the 15 objects referenced by the **leak-import** command is allowed to be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters while the remaining 14 objects have a maximum length of 64 characters each.

When multiple **leak-import** commands are issued, the last command entered overrides the previous command.

When a **leak-import** policy is not specified, no BGP routes from other routing instances are leaked into the VPRN BGP RIB.

The **no** form of the command removes the policy association.

- Parameters** *plcy-or-long-expr* — The route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long). Allowed values are any string up to 255 characters long 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.
- plcy-or-expr* — The route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long). Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

route-table-import

- Syntax** **route-table-import** *policy-name*
no route-table-import
- Context** config>service>vprn>bgp>rib-management>ipv4
config>service>vprn>bgp>rib-management>label-ipv4
config>service>vprn>bgp>rib-management>ipv6
- Description** This command specifies the name of a route policy to control the importation of active routes from the IP route table into one of the BGP RIBs.
- If the **route-table-import** command is not configured, or if the command refers to an empty policy, all non-BGP routes from the IP route table are imported into the applicable RIB.
- If the **route-table-import** command is configured, then routes dropped or rejected by the configured policy are not installed in the associated RIB. Rejected routes cannot be advertised to BGP peers associated with the RIB, but they can still be used to resolve BGP next-hops of routes in that RIB. If the active route for a prefix is rejected by the **route-table-import** policy, then the best BGP route for that prefix in the BGP RIB can be advertised to peers as though it is used.
- Aggregate routes are always imported into each RIB, independent of the **route-table-import** policy.
- Route modifications specified in the actions of a **route-table-import** policy are ignored and have no effect on the imported routes.
- Default** no route-table-import

Parameters *policy-name* — Specifies the name of a policy-statement (up to 64 characters).

split-horizon

Syntax	split-horizon no split-horizon
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command enables the use of split-horizon. When applied globally, to a group, or a specific peer, split-horizon prevents routes from being reflected back to a peer that sends the best route. It applies to routes of all address families and to any type of sending peer; confed-EBGP, EBGp and IBGP. The configuration default is no split-horizon, meaning that no effort is taken to prevent a best route from being reflected back to the sending peer.  Caution: Use of the split-horizon command may have a detrimental impact on peer and route scaling and therefore operators are encouraged to use it only when absolutely needed. The no form of the command disables split horizon command which allows the lower level to inherit the setting from an upper level.
Default	no split-horizon

third-party-nexthop

Syntax	third-party-nexthop no third-party-nexthop
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	Use this command to enable the router to send third-party next-hop to EBGp peers in the same subnet as the source peer, as described in RFC 4271. If enabled when an IPv4 or IPv6 route is received from one EBGp peer and advertised to another EBGp peer in the same IP subnet, the BGP next-hop is left unchanged. Third-party next-hop is not done if the address family of the transport does not match the address family of the route.

The **no** form of the command prevents BGP from performing any third party next-hop processing toward any single-hop EBGP peers within the scope of the command. No third-party next-hop means the next-hop will always carry the IP address of the interface used to establish the TCP connection to the peer.

Default no third-party-nexthop

type

Syntax [no] type {internal | external}

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command designates the BGP peer as type internal or external.

The type of **internal** indicates the peer is an IBGP peer while the type of external indicates that the peer is an EBGP peer.

By default, the OS derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered **internal**. If the local AS is different, then the peer is considered **external**.

The **no** form of the command used at the group level reverts to the default value.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default no type

Parameters **internal** — Configures the peer as internal.
external — Configures the peer as external.
no type — Type of neighbor is derived on the local AS specified.

updated-error-handling

Syntax [no] updated-error-handling

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command controls whether SR OS utilizes the new neighbor-complete bit when processing optional transitive path attributes and advertising them to the associated BGP neighbor.

This command also control if SR OS utilizes the error handling mechanism for optional-transitive path attributes.

Default no updated-error-handling

ttl-security

Syntax **ttl-security** *min-ttl-value*
no ttl-security

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description Configure TTL security parameters for incoming packets.

Parameters *min-ttl-value* — Specifies the minimum TTL value for an incoming BGP packet.

Values 1 to 255

Default 1

3.8.2.29 ETH-CFM Service Commands

eth-cfm

Syntax **eth-cfm**

Context config>service>vprn
config>service>vprn>if>sap
config>service>vprn>if>spoke-sdp
config>service>vprn>sub-if>grp-if>sap

Description This command enters the context to configure ETH-CFM parameters.

collect-lmm-stats

Syntax **collect-lmm-stats**
no collect-lmm-stats

Context config>service>vprn>if>sap>eth-cfm
config>service>vprn>if>spoke-sdp>eth-cfm
config>service>vprn>sub-if>grp-if>sap>eth-cfm

Description This command enables the collection of statistics on the SAP or MPLS SDP binding on which the ETH- LMM test is configured. The collection of LMM statistics must be enabled if a MEP is launching or responding to ETH-LMM packets. If LMM statistics collection is not enabled, the counters in the LMM and LMR PDU do not represent accurate measurements and all measurements should be ignored. The **show sap-using eth-cfm collect-lmm-stats** command and the **show sdp-using eth-cfm collect-lmm-stats** command can be used to display which entities are collecting stats.

The **no** form of the command disables and deletes the counters for this SAP or MPLS SDP binding.

Default no collect-lmm-stats

collect-lmm-fc-stats

Syntax **collect-lmm-fc-stats**

Context config>service>vprn>if>sap>eth-cfm
config>service>vprn>if>spoke-sdp>eth-cfm
config>service>vprn>sub-if>grp-if>sap>eth-cfm

Description This command enters the context to configure per-forwarding class (FC) LMM information collection.

This command is mutually exclusive with the **collect-lmm-stats** command when there is entity resource contention.

fc

Syntax **fc** *fc-name* [*fc-name*]
no fc

Context config>service>vprn>if>sap>eth-cfm>collect-lmm-fc-stats
config>service>vprn>if>spoke-sdp>eth-cfm>collect-lmm-fc-stats
config>service>vprn>sub-if>grp-if>sap>eth-cfm>collect-lmm-fc-stats

Description This command creates individual counters for the specified FCs without regard for profile. All countable packets that match a configured FC, regardless of profile, will be included in this counter.

A differential is performed when this command is re-entered. Omitted FCs will stop counting, newly added FCs will start counting, and unchanged FCs will continue to count.

Up to eight FCs may be specified. An FC that is specified as part of this command for this specific context cannot be specified as a profile-aware FC using the **fc-in-profile** command under the same context.

The **no** form of the command removes all previously defined FCs and stops counting for those FCs.

Default	no fc
Parameters	<i>fc-name</i> — Specifies the name of the FC for which to create an individual profile-unaware counter. Up to 8 FCs can be named in a single statement. In order for the counter to be used, the config>oam-pm>session>ethernet>priority command must be configured with a numerical value representing the FC name (7 = NC, 6 = H1, 5 = EF, 4 = H2, 3 = L1, 2 = AF, 1 = L2, 0 = BE), and the config>oam-pm>session>ethernet>lmm>enable-fc-collection command must be enabled. Values nc, h1, ef, h2, l1, af, l2, be

fc-in-profile

Syntax	fc-in-profile <i>fc-name</i> [<i>fc-name</i>] no fc-in-profile
Context	config>service>vprn>if>sap>eth-cfm>collect-lmm-fc-stats config>service>vprn>if>spoke-sdp>eth-cfm>collect-lmm-fc-stats config>service>vprn>sub-if>grp-if>sap>eth-cfm>collect-lmm-fc-stats
Description	This command creates individual counters for the specified FCs with regard for profile. All countable packets that match a configured FC and are deemed to be in profile will be included in this counter. A differential is performed when this command is re-entered. Omitted FCs will stop counting, newly added FCs will start counting, and unchanged FCs will continue to count. Up to eight FCs may be specified. An FC that is specified as part of this command for this specific context cannot be specified as a profile-unaware FC using the fc command under the same context. The no form of the command removes all previously defined FCs and stops counting for those FCs.
Default	no fc-in-profile
Parameters	<i>fc-name</i> — Specifies the name of the FC for which to create an individual profile-aware counter. Up to 8 FCs can be named in a single statement. In order for the counter to be used, the config>oam-pm>session>ethernet>priority command must be configured with a numerical value representing the FC name (7 = NC, 6 = H1, 5 = EF, 4 = H2, 3 = L1, 2 = AF, 1 = L2, 0 = BE), and the config>oam-pm>session>ethernet>lmm>enable-fc-collection command must be enabled. Values nc, h1, ef, h2, l1, af, l2, be

mep

Syntax	mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [direction { up down }] no mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i>
Context	config>service>vprn>if>sap>eth-cfm config>service>vprn>if>spoke-sdp>eth-cfm config>service>vprn>sub-if>grp-if>sap>eth-cfm
Description	This command configures the ETH-CFM maintenance endpoint (MEP).
Parameters	<i>mep-id</i> — Specifies the maintenance association end point identifier. Values 1 to 8191 <i>md-index</i> — Specifies the maintenance domain (MD) index value. Values 1 to 4294967295 <i>ma-index</i> — Specifies the MA index value. Values 1 to 4294967295 direction up down — Indicates the direction in which the maintenance association (MEP) faces on the bridge port. Direction UP is not supported on VPRN MEPs. Values down — Sends continuity check messages away from the MAC relay entity. up — Sends continuity check messages towards the MAC relay entity.

ais-enable

Syntax	[no] ais-enable
Context	config>service>vprn>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm config>service>vprn>sub-if>grp-if>sap>eth-cfm
Description	This command configures the reception of Alarm Indication Signal (AIS) message.

interface-support-enable

Syntax	[no] interface-support-enable
Context	config>service>vprn>sap>eth-cfm>mep>ais-enable config>service>vprn>spoke-sdp>eth-cfm>mep>ais-enable

Description This command enables the AIS function to consider the operational state of the entity on which it is configured. With this command, ETH-AIS on DOWN MEPs will be triggered and cleared based on the operational status of the entity on which it is configured. If CCM is also enabled then transmission of the AIS PDU will be based on either the non-operational state of the entity or on ANY CCM defect condition. AIS generation will cease if BOTH operational state is UP and CCM has no defect conditions. If the MEP is not CCM enabled then the operational state of the entity is the only consideration assuming this command is present for the MEP.

The no form of this command means that AIS will not be generated or stopped based on the state of the entity on which the DOWN MEP is configured.

Default no interface-support-enable

ccm-enable

Syntax [no] ccm-enable

Context config>service>vprn>if>sap>eth-cfm>mep
config>service>vprn>if>spoke-sdp>eth-cfm>mep
config>service>vprn>sub-if>grp-if>sap>eth-cfm

Description This command enables the generation of CCM messages.

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

ccm-ltm-priority

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

Context config>service>vprn>if>sap>eth-cfm>mep
config>service>vprn>if>spoke-sdp>eth-cfm>mep
config>service>vprn>sub-if>grp-if>sap>eth-cfm

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

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

Default The highest priority on the bridge-port.

Parameters *priority* — Specifies the priority of CCM and LTM messages.

Values 0 to 7

ccm-padding-size

Syntax	[no] ccm-padding-size <i>ccm-padding</i>
Context	config>service>vprn>if>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm>mep config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep
Description	This command sets the byte size of the optional Data TLV to be included in the ETH-CC PDU. This will increase the size of the ETH-CC PDU by the configured value. The base size of the ETH-CC PDU, including the Interface Status TLV and Port Status TLV, is 83 bytes not including the Layer Two encapsulation. CCM padding is not supported when the CCM-Interval is less than one second.
Parameters	<i>ccm-padding</i> — Specifies the byte size of the Optional Data TLV.
Values	3 to 1500

csf-enable

Syntax	[no] csf-enable
Context	config>service>vprn>if>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm>mep config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep
Description	This command enables the reception and local processing of ETH-CSF frames.

multiplier

Syntax	multiplier <i>multiplier-value</i> no multiplier
Context	config>service>vprn>if>sap>eth-cfm>mep>csf-enable config>service>vprn>if>spoke-sdp>eth-cfm>mep>csf-enable config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep>csf-enable
Description	This command enables the multiplication factor applied to the receive time used to clear the CSF condition in increments of 5.
Default	multiplier 3.5
Parameters	<i>multiplier-value</i> — Specifies the multiplier used for timing out CSF.
Values	0.0, 2.0 to 30.0

eth-test-enable

Syntax	[no] eth-test-enable
Context	config>service>vprn>if>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm>mep config>service>vprn>sub-if>grp-if>sap>eth-cfm
Description	This command enables eth-test functionality on MEP. For this test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands: <pre>oam eth-cfm eth-test <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [priority <i>priority</i>] [data-length <i>data-length</i>]</pre> A check is done for both the provisioning and test to ensure the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the CLI and SNMP will indicate the problem.

test-pattern

Syntax	test-pattern {all-zeros all-ones} [crc-enable] no test-pattern
Context	config>service>vprn>if>sap>eth-cfm>mep>eth-test-enable config>service>vprn>if>spoke-sdp>eth-cfm>mep>eth-test-enable config>service>vprn>sub-if>grp-if>sap>eth-cfm>eth-test-enable
Description	This command configures the test pattern for eth-test frames. The no form of the command removes the values from the configuration.
Default	test-pattern all-zeros
Parameters	all-zeros — Specifies to use all zeros in the test pattern. all-ones — Specifies to use all ones in the test pattern. crc-enable — generates a CRC checksum.

bit-error-threshold

Syntax	bit-error-threshold <i>bit-errors</i>
Context	config>service>vprn>if>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm>mep config>service>vprn>sub-if>grp-if>sap>eth-cfm
Description	This command specifies the lowest priority defect that is allowed to generate a fault alarm.

Default	bit-error-threshold 1
Parameters	<i>bit-errors</i> — Specifies the lowest priority defect.
Values	0 to 11840

one-way-delay-threshold

Syntax	one-way-delay-threshold <i>time</i>
Context	config>service>vprn>if>sap>eth-cfm config>service>vprn>if>spoke-sdp>eth-cfm config>service>vprn>sub-if>grp-if>sap>eth-cfm
Description	This command enables one way delay threshold time limit.
Default	3 seconds
Parameters	<i>priority</i> — Specifies the value for the threshold.
Values	0 to 600

squelch-ingress-levels

Syntax	squelch-ingress-levels [<i>md-level</i> [<i>md-level</i> ...]] no squelch-ingress-levels
Context	config>service>vprn>if>sap>eth-cfm config>service>vprn>if>spoke-sdp>eth-cfm config>service>vprn>sub-if>grp-if>sap>eth-cfm
Description	This command defines the levels of the ETH-CFM PDUs that will silently be discarded on ingress into the SAP or SDP Binding from the wire. All ETH-CFM PDUs inbound to the SAP or SDP binding will be dropped that match the configured levels without regard for any other ETH-CFM criteria. No statistical information or drop count will be available for any ETH-PDU that is silently discarded by this option. The operator must configure a complete contiguous list of md-levels up to the highest level that will be dropped. The command must be retyped in complete form to modify a previous configuration, if the operator does not want to delete it first. The no form of the command removes the silent discarding of previously matching ETH-CFM PDUs.
Default	no squelch-ingress-levels
Parameters	<i>md-level</i> — Identifies the level.
Values	[0 to 7]

tunnel-fault

Syntax	tunnel-fault {accept ignore}
Context	config>service>vprn>eth-cfm config>service>vprn>if>sap>eth-cfm config>service>vprn>sub-if>grp-if>sap>eth-cfm
Description	Allows the individual service SAPs to react to changes in the tunnel MEP state. When tunnel-fault accept is configured at the service level, the SAP will react according to the service type, Epipe will set the operational flag and VPLS, IES and VPRN SAP operational state will become down on failure or up on clear. This command triggers the OAM mapping functions to mate SAPs and bindings in an Epipe service as well as setting the operational flag. If AIS generation is the requirement for the Epipe services this command is not required. See the ais-enable command in the epipe>sap>eth-cfm>ais-enable context for more information. This works in conjunction with the tunnel-fault accept on the individual SAPs. Both must be set to accept to react to the tunnel MEP state. By default the service level command is “ignore” and the sap level command is “accept”. This means simply changing the service level command to “accept” will enable the feature for all SAPs. This is not required for Epipe services that only wish to generate AIS on failure.
Default	tunnel-fault ignore (Service Level) tunnel-fault accept (SAP Level for Epipe and VPLS)
Parameters	<i>accept</i> — Share fate with the facility tunnel MEP. <i>ignore</i> — Do not share fate with the facility tunnel MEP.

fault-propagation-enable

Syntax	fault-propagation-enable {use-if-tlv suspend-ccm} no fault-propagation-enable
Context	config>service>vprn>if>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm>mep config>service>vprn>sub-if>grp-if>sap>eth-cfm
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.

grace

Syntax	grace
Context	config>service>vprn>if>sap>eth-cfm>mep

```
config>service>vprn>if>spoke-sdp>eth-cfm>mep
config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep
```

Description This command enters the context to configure Nokia ETH-CFM Grace and ITU-T Y.1731 ETH-ED expected defect functional parameters.

eth-ed

Syntax **eth-ed**

Context config>service>vprn>if>sap>eth-cfm>mep>grace
config>service>vprn>if>spoke-sdp>eth-cfm>mep>grace
config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep>grace

Description This command enters the context to configure ITU-T Y.1731 ETH-ED expected defect functional parameters.

max-rx-defect-window

Syntax **max-rx-defect-window** *seconds*
no max-rx-defect-window

Context config>service>vprn>if>sap>eth-cfm>mep>grace>eth-ed
config>service>vprn>if>spoke-sdp>eth-cfm>mep>grace>eth-ed
config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep>grace>eth-ed

Description This command limits the duration of the received ETH-ED expected defect window to the lower value of either the received value from the peer or this parameter.

The **no** form of the command removes the limitation, and any valid defect window value received from a peer MEP in the ETH-ED PDU will be used.

Default no max-rx-defect-window

Parameters *seconds* — Specifies the duration, in seconds, of the maximum expected defect window

Values 1 to 86400

priority

Syntax **priority** *priority*
no priority

Context config>service>vprn>if>sap>eth-cfm>mep>grace>eth-ed
config>service>vprn>if>spoke-sdp>eth-cfm>mep>grace>eth-ed
config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep>grace>eth-ed

Description	This command sets the priority bits and determines the forwarding class based on the mapping of priority to FC. The no form of the command disables the local priority configuration and sets the priority to the ccm-ltm-priority associated with this MEP.
Default	no priority
Parameters	<i>priority</i> — Specifies the priority bit.
Values	0 to 7

rx-eth-ed

Syntax	[no] rx-eth-ed
Context	<pre>config>service>vprn>if>sap>eth-cfm>mep>grace>eth-ed config>service>vprn>if>spoke-sdp>eth-cfm>mep>grace>eth-ed config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep>grace>eth-ed</pre>
Description	This command enables the reception and processing of the ITU-T Y.1731 ETH-ED PDU on the MEP. The no form of the command disables the reception of the ITU-T Y.1731 ETH-ED PDU on the MEP.
Default	rx-eth-ed

tx-eth-ed

Syntax	[no] tx-eth-ed
Context	<pre>config>service>vprn>if>sap>eth-cfm>mep>grace>eth-ed config>service>vprn>if>spoke-sdp>eth-cfm>mep>grace>eth-ed config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep>grace>eth-ed</pre>
Description	This command enables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP when a system soft reset notification is received for one or more cards. The config>eth-cfm>system>grace-tx-enable command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted. The no form of the command disables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP.
Default	no tx-eth-ed

eth-vsm-grace

Syntax	eth-vsm-grace
Context	config>service>vprn>if>sap>eth-cfm>mep>grace config>service>vprn>if>spoke-sdp>eth-cfm>mep>grace config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep>grace
Description	This command enters the context to configure Nokia ETH-CFM Grace functional parameters.

rx-eth-vsm-grace

Syntax	[no] rx-eth-vsm-grace
Context	config>service>vprn>if>sap>eth-cfm>mep>grace>eth-vsm-grace config>service>vprn>if>spoke-sdp>eth-cfm>mep>grace>eth-vsm-grace config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep>grace>eth-vsm-grace
Description	This command enables the reception and processing of the Nokia ETH-CFM Grace PDU on the MEP. The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect. The no form of the command disables the reception of the Nokia ETH-CFM Grace PDU on the MEP.
Default	rx-eth-vsm-grace

tx-eth-vsm-grace

Syntax	[no] tx-eth-vsm-grace
Context	config>service>vprn>if>sap>eth-cfm>mep>grace>eth-vsm-grace config>service>vprn>if>spoke-sdp>eth-cfm>mep>grace>eth-vsm-grace config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep>grace>eth-vsm-grace
Description	This command enables the transmission of the Nokia ETH-CFM Grace PDU from the MEP when a system soft reset notification is received for one or more cards. The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect. The config>eth-cfm>system>grace-tx-enable command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted.

The **no** form of the command disables the transmission of the Nokia ETH-CFM Grace PDU from the MEP.

Default tx-eth-vsm-grace

low-priority-defect

Syntax	low-priority-defect {allDef macRemErrXcon remErrXcon errXcon xcon noXcon}
Context	config>service>vprn>if>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm>mep config>service>vprn>sub-if>grp-if>sap>eth-cfm
Description	This command specifies the lowest priority defect that is allowed to generate a fault alarm.
Default	low-priority-defect macRemErrXcon
Parameters	<i>parameters</i> — Specifies the lowest priority defect.
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

3.8.2.30 IS-IS Commands

isis

Syntax	[no] isis <i>isis-instance</i>
Context	config>service>vprn
Description	This command creates the context to configure the Intermediate-System-to-Intermediate-System (IS-IS) protocol instance in the VPRN. The IS-IS protocol instance is enabled with the no shutdown command in the config>service>vprn>isis context. Alternatively, the IS-IS protocol instance is disabled with the shutdown command in the config>service>vprn>isis context.

IS-IS instances are **shutdown** when created, so that all parameters can be configured prior to the instance being enabled.

The **no** form of the command disables the ISIS protocol instance from the given VPRN service.

Default	0
Parameters	<i>isis-instance</i> — Specifies the instance ID for an IS-IS instance.
Values	1 to 31

advertise-passive-only

Syntax	[no] advertise-passive-only
Context	config>service>vprn>isis
Description	This command enables and disables IS-IS for the VPRN instance to advertise only prefixes that belong to passive interfaces.

advertise-router-capability

Syntax	advertise-router-capability {area as} no advertise-router-capability
Context	config>service>vprn>isis config>service>vprn>isis>level
Description	This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A new TLV as defined in RFC 4971 advertises the TE Node Capability Descriptor capability. The parameters (area & as) control the scope of the capabilities advertisements. The no form of this command, disables this capability.
Default	no advertise-router-capability
Parameters	area — Capabilities are only advertised within the area of origin. as — Capabilities are only advertised throughout the entire autonomous system.

all-l1isis

Syntax	all-l1isis ieee-address no all-l1isis

Context	config>service>vprn>isis
Description	This command specifies the MAC address to use for the VPRN instance of the L1 IS-IS routers. The MAC address should be a multicast address.
Default	all-l1isis 01:80:c2:00:00:14
Parameters	<i>ieee-address</i> — Specifies the destination MAC address for all L1 I-IS neighbors on the link for this ISIS instance.

all-l2isis

Syntax	all-l2isis <i>ieee-address</i> no all-l2isis
Context	config>service>vprn>isis
Description	This command specifies the MAC address to use for L2 IS-IS routers for the VPRN instance. The MAC address should be a multicast address.
Default	all-l2isis 01:80:c2:00:00:15
Parameters	<i>ieee-address</i> — Specifies the destination MAC address for all L2 ISIS neighbors on the link for this ISIS instance.

area-id

Syntax	[no] area-id <i>area-address</i>
Context	config>service>vprn>isis
Description	This command configures the area ID portion of NSAP addresses for the VPRN instance. This identifies a point of connection to the network, such as a router interface, and is called a Network Service Access Point (NSAP). Addresses in the IS-IS protocol are based on the ISO NSAP addresses and Network Entity Titles (NETs), not IP addresses.

A maximum of 3 **area addresses** can be configured for the VPRN instance.

NSAP addresses are divided into three parts. Only the area ID portion is configurable.

- Area ID — A variable length field between 1 and 13 bytes long. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.
- System ID — A six-byte system identification. This value is not configurable. The system ID is derived from the system or router ID.
- Selector ID — A one-byte selector identification that must contain zeros when configuring a NET. This value is not configurable. The selector ID is always 00.

The NET is constructed like an NSAP but the selector byte contains a 00 value. NET addresses are exchanged in hello and LSP PDUs. All net addresses configured on the node are advertised to its neighbors.

For Level 1 interfaces, neighbors can have different area IDs, but, they must have at least one area ID (AFI + area) in common. Sharing a common area ID, they become neighbors and area merging between the potentially different areas can occur.

For Level 2 (only) interfaces, neighbors can have different area IDs. However, if they have no area IDs in common, they become only Level 2 neighbors and Level 2 LSPs are exchanged.

For Level 1 and Level 2 interfaces, neighbors can have different area IDs. If they have at least one area ID (AFI + area) in common, they become neighbors. In addition to exchanging Level 2 LSPs, area merging between potentially different areas can occur.

If multiple **area-id** commands are entered, the system ID of all subsequent entries must match the first **area** address.

The **no** form of the command removes the area address.

auth-keychain

Syntax	auth-keychain <i>name</i>
Context	config>service>vprn>isis< config>service>vprn>isis>level
Description	This command configures an authentication keychain to use for the protocol interface for the VPRN instance. The keychain allows the rollover of authentication keys during the lifetime of a session.
Default	no auth-keychain
Parameters	<i>name</i> — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

authentication-check

Syntax	[no] authentication-check
Context	config>service>vprn>isis
Description	This command sets an authentication check to reject PDUs that do not match the type or key requirements for the VPRN instance. The default behavior when authentication is configured is to reject all IS-IS protocol PDUs that have a mismatch in either the authentication type or authentication key.

When **no authentication-check** is configured, authentication PDUs are generated and IS-IS PDUs are authenticated on receipt. However, mismatches cause an event to be generated and will not be rejected.

The **no** form of this command allows authentication mismatches to be accepted and generate a log event.

Default authentication-check — Rejects authentication mismatches.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>service>vprn>isis
config>service>vprn>isis>level

Description This command sets the authentication key used to verify PDUs sent by neighboring routers on the interface for the VPRN instance.

Neighboring routers use passwords to authenticate PDUs sent from an interface. For authentication to work, both the authentication *key* and the authentication *type* on a segment must match. The [OSPF Commands](#) statement must also be included.

To configure authentication on the global level, configure this command in the **config>router>isis** context. When this parameter is configured on the global level, all PDUs are authenticated including the hello PDU.

To override the global setting for a specific level, configure the **authentication-key** command in the **config>router>isis>level** context. When configured within the specific level, hello PDUs are not authenticated.

The **no** form of the command removes the authentication key.

Default no authentication-key — No authentication key is configured.

Parameters *authentication-key* — The authentication key. The key can be any combination of ASCII characters up to 255 characters in length (un-encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key — The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

authentication-type

Syntax	authentication-type {password message-digest} no authentication
Context	config>service>vprn>isis config>service>vprn>isis>level
Description	This command enables either simple password or message digest authentication or must go in either the global IS-IS or IS-IS level context. Both the authentication key and the authentication type on a segment must match. The authentication-key statement must also be included. Configure the authentication type on the global level in the config>router>isis context. Configure or override the global setting by configuring the authentication type in the config>router>isis>level context. The no form of the command disables authentication.
Default	no authentication-type — No authentication type is configured and authentication is disabled.
Parameters	password — Specifies that simple password (plain text) authentication is required. message-digest — Specifies that MD5 authentication in accordance with RFC2104 is required.

csnp-authentication

Syntax	[no] csnp-authentication
Context	config>service>vprn>isis config>service>vprn>isis>level
Description	This command enables authentication of individual ISIS packets of complete sequence number PDUs (CSNP) type for the VPRN instance.

default-route-tag

Syntax	default-route-tag tag

no default-route-tag

Context	config>service>vprn>isis
Description	This command configures the route tag for default route for the router or VPRN service.
Parameters	<i>tag</i> — Assigns a default tag.
Values	1 — 4294967295

export

Syntax	[no] export <i>policy-name</i> [<i>policy-name</i> ...up to 5 max]
Context	config>service>vprn>isis
Description	This command configures export routing policies that determine the routes exported from the routing table to IS-IS. If no export policy is defined, non IS-IS routes are not exported from the routing table manager to IS-IS. If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered overrides the previous command. A maximum of five policy names can be specified. If an aggregate command is also configured in the config>router context, then the aggregation is applied before the export policy is applied. Routing policies are created in the config>router>policy-options context. The no form of the command removes the specified <i>policy-name</i> or all policies from the configuration if no <i>policy-name</i> is specified.
Default	no export — No export policy name is specified.
Parameters	<i>policy-name</i> — The export policy name. Up to five <i>policy-name</i> arguments can be specified.

export-limit

Syntax	export-limit <i>number</i> [<i>log percentage</i>] no export-limit
Context	config>service>vprn>isis
Description	This command configures the maximum number of routes (prefixes) that can be exported into IS-IS from the route table for the VPRN instance.

The **no** form of the command removes the parameters from the configuration.

Default	no export-limit - The export limit for routes or prefixes is disabled.
Parameters	<i>number</i> — Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table. Values 1 to 4294967295 <i>log percentage</i> — Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent. Values 1 to 100

graceful-restart

Syntax	[no] graceful-restart
Context	config>service>vprn>isis
Description	This command enables IS-IS graceful restart (GR) to minimize service interruption. When the control plane of a GR-capable router fails or restarts, the neighboring routers (GR helpers) temporarily preserve IS-IS forwarding information. Traffic continues to be forwarded to the restarting router using the last known forwarding tables. If the control plane of the restarting router becomes operationally and administratively up within the grace period, the restarting router resumes normal IS-IS operation. If the grace period expires, then the restarting router is presumed inactive and the IS-IS topology is recalculated to route traffic around the failure. The no form of the command disables graceful restart and removes the graceful restart configuration from the IS-IS instance. Default no graceful-restart

helper-disable

Syntax	[no] helper-disable
Context	config>service>vprn>isis>graceful-restart
Description	This command disables helper support for IS-IS graceful restart (GR). When graceful-restart is enabled, the router can be a helper (meaning that the router is helping a neighbor to restart), a restarting router, or both. The router only supports helper mode. It will not act as a restarting router, because the high availability feature set already preserves IS-IS forwarding information such that this functionality is not needed. This command is a historical command and should not be disabled. Configuring helper-disable has the effect of disabling graceful restart, because the router only supports helper mode. The no helper-disable command enables helper support and is the default when graceful restart is enabled.

Default no helper-disable

hello-authentication

Syntax [no] hello-authentication

Context config>service>vprn>isis
config>service>vprn>isis>interface
config>service>vprn>isis>level

Description This command enables authentication of individual IS-IS Hello packets for the VPRN instance.

The **no** form of the command suppresses authentication of Hello packets.

hello-padding

Syntax hello-padding {none | adaptive | loose | strict}
no hello-padding

Context config>service>vprn>isis
config>service>vprn>isis>level
config>service>vprn>isis>interface
config>service>vprn>isis>interface>level

Description This command enables the IS-IS Hello (IIH) message padding to ensure that IS-IS LSPs can traverse the link. When this option is enabled, IS-IS Hello messages are padded to the maximum LSP MTU value, which can be set with the **lsp-mtu-size** command.

The **no** form of the command disables IS-IS Hello message padding at this level. However, the router may still perform hello padding if it was set at a higher level in the configuration. To ensure that Hello message padding is disabled, set all levels of configuration to **no hello-padding**.

Default no hello-padding

Parameters **adaptive** — Specifies the adaptive padding option; this option is able to detect MTU asymmetry from one side of the connection but uses more overhead than loose padding.

- point-to-point interface—Hello PDUs are padded until the sender declares an adjacency on the link to be in the state up. If the implementation supports RFC 3373/5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*, then this is when the three-way state is up. If the implementation uses the “classic” algorithm described in ISO 10589, this is when the adjacency state is up. If the neighbor does not support the adjacency state TLV, then padding continues.
- broadcast interface—Padding starts until at least one adjacency is up on the interface.

loose — Specifies the loose padding option; the loose padding may not be able to detect certain conditions such as asymmetrical MTUs between the routing devices.

- point-to-point interface—the Hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the INIT state
- broadcast interface—padding starts until at least one adjacency (broadcast only has up/down) is up on the interface

none — Specifies that the Hello message padding is not enabled at this level, even if it is configured at one of the parent levels.

strict — Specifies the strict padding option.

- point-to-point interface—padding is done for all adjacency states, and is continuous. Strict padding has the most overhead but detects MTU issues on both sides of a link
- broadcast interface—padding is done for all adjacency states, and is continuous. Strict padding has the most overhead but detects MTU issues on both sides of a link

ignore-lsp-errors

Syntax	[no] ignore-lsp-errors
Context	config>service>vprn>isis
Description	This command specifies that for this VPRN instance, ISIS will ignore LSP packets with errors. When enabled, IS-IS LSP errors will be ignored and the associated record will not be purged. This command enables ISIS to ignore the ATT bit and therefore suppress the installation of default routes. The no form of the command specifies that ISIS will not ignore LSP errors.

iid-tlv-enable

Syntax	[no] iid-tlv-enable
Context	config>service>vprn>isis
Description	This command enables IS-IS multi-instance (MI) as described in draft-ietf-isis-mi-02. Multiple instances allow instance-specific adjacencies to be formed that support multiple network topologies on the same physical interfaces. Each instance has an LSDB, and each PDU contains a TLV identifying the instance and the topology to which the PDU belongs. The iid-tlv-enable (based on draft-ietf-isis-mi-02) and standard-multi-instance (based on draft-ginsberg-isis-mi-bis-01) commands cannot be configured in the same instance, because the MAC addresses and PDUs in each standard are incompatible.
Default	no iid-tlv-enable

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>service>vprn>isis
Description	This command creates the context to configure an IS-IS interface. When an area is defined, the interfaces belong to that area. Interfaces cannot belong to separate areas. When the interface is a POS channel, the OSINCP is enabled when the interface is created and removed when the interface is deleted. The no form of the command removes IS-IS from the interface. The shutdown command in the config>router>isis>interface context administratively disables IS-IS on the interface without affecting the IS-IS configuration.
Default	no interface — No IS-IS interfaces are defined.
Parameters	<i>ip-int-name</i> — Identify the IP interface name created in the config>router>interface context. The IP interface name must already exist.

bfd-enable

Syntax	[no] bfd-enable {ipv4 ipv6} [include-bfd-tlv]
Context	config>service>vprn>isis>interface
Description	This command enables the use of bi-directional forwarding (BFD) to control IPv4 adjacencies. By enabling BFD on an IPv4 or IPv6 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 by the BFD command under the IP interface. This command must be given separately to enable/disable BFD for both IPv4 and IPv6. The no form of this command removes BFD from the associated adjacency.
Default	no bfd-enable ipv4 no bfd-enable ipv6

csnp-interval

Syntax	csnp-interval <i>seconds</i> no csnp-interval
Context	config>service>vprn>isis>interface

Description	This command configures the time interval, in seconds, to send complete sequence number (CSN) PDUs from the interface. IS-IS must send CSN PDUs periodically. The no form of the command reverts to the default value.
Default	csnp-interval 10 — CSN PDUs are sent every 10 seconds for LAN interfaces. csnp-interval 5 — CSN PDUs are sent every 5 seconds for point-to-point interfaces.
Parameters	<i>seconds</i> — The time interval, in seconds between successive CSN PDUs sent from this interface expressed as a decimal integer. 1 to 65535

default-instance

Syntax	[no] default-instance
Context	config>service>vprn>isis>interface
Description	This command enables a non-MI capable router to establish an adjacency and operate with an SR OS in a non-zero instance. If the router does not receive IID-TLVs, it will establish an adjacency in a single instance. Instead of establishing an adjacency in the standard instance 0, the router will establish an adjacency in the configured non-zero instance. The router will then operate in the configured non-zero instance so that it appears to be in the standard instance 0 to its neighbor. This feature is supported on point-to-point interfaces, broadcast interfaces are not supported. The no form of the command disables the functionality so that the router can only establish adjacencies in the standard instance 0.
Default	no default-instance

hello-auth-keychain

Syntax	hello-auth-keychain <i>name</i>
Context	config>service>vprn>isis>interface config>service>vprn>isis>interface>level
Description	This command configures an authentication keychain to use for the protocol interface. The keychain allows the rollover of authentication keys during the lifetime of a session.
Default	no hello-auth-keychain
Parameters	<i>name</i> — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

hello-authentication-key

Syntax	hello-authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no hello-authentication-key
Context	config>service>vprn>isis>interface config>service>vprn>isis>interface>level
Description	This command configures the authentication key (password) for hello PDUs. Neighboring routers use the password to verify the authenticity of hello PDUs sent from this interface. Both the hello authentication key and the hello authentication type on a segment must match. The hello-authentication-type must be specified. To configure the hello authentication key in the interface context use the hello-authentication-key in the config>router>isis>interface context. To configure or override the hello authentication key for a specific level, configure the hello-authentication-key in the config>router>isis>interface>level context. If both IS-IS and hello-authentication are configured, Hello messages are validated using hello authentication. If only IS-IS authentication is configured, it will be used to authenticate all IS-IS (including hello) protocol PDUs. When the hello authentication key is configured in the config>router>isis>interface context, it applies to all levels configured for the interface. The no form of the command removes the authentication-key from the configuration.
Default	no hello-authentication-key — No hello authentication key is configured.
Parameters	<i>authentication-key</i> — The hello authentication key (password). The key can be any combination of ASCII characters up to 254 characters in length (un-encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" "). <i>hash-key</i> — The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" "). This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided. hash — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified. hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the hash2 encrypted variable cannot be copied and pasted. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.

hello-authentication-type

Syntax	hello-authentication-type {password message-digest} no hello-authentication-type
Context	config>service>vprn>isis>interface config>service>vprn>isis>interface>level
Description	This command enables hello authentication at either the interface or level context. Both the hello authentication key and the hello authentication type on a segment must match. The hello authentication-key statement must also be included. To configure the hello authentication type at the interface context, use hello-authentication-type in the config>router>isis>interface context. To configure or override the hello authentication setting for a given level, configure the hello-authentication-type in the config>router>isis>interface>level context. The no form of the command disables hello authentication.
Default	no hello-authentication-type — hello authentication is disabled
Parameters	password — Specifies simple password (plain text) authentication is required. message-digest — Specifies MD5 authentication in accordance with RFC2104 (HMAC: Keyed-Hashing for Message Authentication) is required.

interface-type

Syntax	interface-type {broadcast point-to-point} no interface-type
Context	config>service>vprn>isis>interface
Description	This command configures the IS-IS 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 designated IS-IS overhead if the link is used as a point-to-point. If the interface type is not known at the time the interface is added to IS-IS and subsequently the IP interface is bound (or moved) to a different interface type, then this command must be entered manually. The no form of the command reverts to the default value.
Default	point-to-point — For IP interfaces on SONET channels. broadcast — For IP interfaces on Ethernet or unknown type physical interfaces.
Special Cases	SONET — Interfaces on SONET channels default to the point-to-point type.

Ethernet or Unknown — Physical interfaces that are Ethernet or unknown default to the broadcast type.

Parameters **broadcast** — Configures the interface to maintain this link as a broadcast network.
point-to-point — Configures the interface to maintain this link as a point-to-point link.

ipv4-multicast-disable

Syntax **[no] ipv4-multicast-disable**

Context config>service>vprn>isis>interface

Description This command administratively disables/enables ISIS operation for IPv4.

Default no ipv4-multicast-disable

ipv6-unicast-disable

Syntax **[no] ipv6-unicast-disable**

Context config>router>isis>interface
config>service>vprn>isis>interface

Description This command disables IS-IS IPv6 unicast routing for the interface.

By default IPv6 unicast on all interfaces is enabled. However, IPv6 unicast routing on IS-IS is in effect when the **config>router>isis>ipv6-routing mt** command is configured.

The **no** form of the command enables IS-IS IPv6 unicast routing for the interface.

hello-interval

Syntax **hello-interval** *seconds*
no hello-interval

Context config>router>isis>interface>level
config>service>vprn>isis>interface>level

Description This command configures the interval in seconds between Hello messages sent on the interface at this level.



Note: The neighbor hold-time is (hello multiplier * hello interval) on point-to-point interfaces, and (hello multiplier * hello interval / 3) on broadcast interfaces. Hello values can be adjusted for faster convergence, but the hold-time should always be > 3 to reduce routing instability.

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

Default	hello-interval 9
Parameters	<i>seconds</i> — The hello interval in seconds expressed as a decimal integer.
Values	1 to 20000

hello-multiplier

Syntax	hello-multiplier <i>multiplier</i> no hello-multiplier
Context	config>router>isis>if>level level-number config>service>vprn>isis>interface>level
Description	This command configures the number of missing hello messages from a neighbor before the router declares the adjacency down.



Note: The neighbor hold-time is (hello multiplier * hello interval) on point-to-point interfaces, and (hello multiplier * hello interval / 3) on broadcast interfaces. Hello values can be adjusted for faster convergence, but the hold-time should always be > 3 to reduce routing instability.

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

Default	hello-multiplier 3
Parameters	<i>multiplier</i> — The multiplier for the hello interval expressed as a decimal integer.
Values	2 to 100

ipv4-multicast-metric

Syntax	ipv4-multicast-metric <i>metric</i> no ipv4-multicast-metric
Context	config>service>vprn>isis>interface>level
Description	This command configures IS-IS interface metric for IPv4 multicast for the VPRN instance. The no form of this command removes the metric from the configuration.
Parameters	<i>metric</i> — Specifies the IS-IS interface metric for IPv4 multicast.
Values	1 to 16777215

ipv6-unicast-metric

Syntax	ipv6-unicast-metric <i>metric</i> no ipv6-unicast-metric
Context	config>service>vprn>isis>interface>level
Description	This command configures IS-IS interface metric for IPv6 unicast. The no form of this command removes the metric from the configuration.
Parameters	<i>metric</i> — Specifies the IS-IS interface metric for IPv6 unicast. Values 1 to 16777215

metric

Syntax	metric <i>metric</i> no metric
Context	config>service>vprn>isis>interface>level
Description	This command configures the metric used for the level on the interface. In order to calculate the lowest cost to reach a given destination, each configured level on each interface must have a cost. The costs for each level on an interface may be different. If the metric is not configured, the default of 10 is used unless reference bandwidth is configured. The no form of the command reverts to the default value.
Default	no metric
Parameters	<i>metric</i> — The metric assigned for this level on this interface. Values 1 to 16777215 Default 10

passive

Syntax	[no] passive
Context	config>service>vprn>isis>interface config>service>vprn>isis>interface>level
Description	This command adds the passive attribute which causes the interface to be advertised as an IS-IS interface without running the IS-IS protocol. Normally, only interface addresses that are configured for IS-IS are advertised as IS-IS interfaces at the level that they are configured.

When the passive mode is enabled, the interface or the interface at the level ignores ingress IS-IS protocol PDUs and will not transmit IS-IS protocol PDUs.

The **no** form of the command removes the passive attribute.

Default passive (service interfaces defined using the service-prefix command in **config>router**)
 no passive (all other interfaces)

priority

Syntax **priority** *number*
 no priority

Context config>service>vprn>isis>interface>level

Description This command configures the priority of the IS-IS router interface for designated router election on a multi-access network.

 This priority is included in hello PDUs transmitted by the interface on a multi-access network. The router with the highest priority is the preferred designated router. The designated router is responsible for sending LSPs with regard to this network and the routers that are attached to it.

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

Default priority 64

Parameters *number* — Specifies the priority for this interface at this level.

Values 0 to 127

sd-offset

Syntax **sd-offset** *offset-value*
 no sd-offset

Context config>service>vprn>isis>interface>level

Description If the pre-FEC error rate of the associated DWDM port crosses the configured **sd-threshold**, this offset-value is added to the IS-IS interface metric. This parameter is only effective if the interface is associated with a DWDM port and the **sd-threshold** value is configured under that port.

 The **no** form of the command reverts the offset value to 0.

Default no sd-offset

Parameters	<i>offset-value</i> — Specifies the amount the interface metric is increased by if the sd-threshold is crossed.
Values	0 to 16777215

sf-offset

Syntax	sf-offset <i>offset-value</i> no sf-offset
Context	config>service>vprn>isis>interface>level
Description	If the pre-FEC error rate of the associated DWDM port crosses the configured sf-threshold, this offset-value is added to the IS-IS interface metric. This parameter is only effective if the interface is associated with a DWDM port and the sf-threshold value is configured under that port. The no form of the command reverts the offset value to 0.
Default	no sf-offset
Parameters	<i>offset-value</i> — Specifies the amount the interface metric is increased by if the sf-threshold is crossed.
Values	0 to 16777215

lfa-policy-map

Syntax	lfa-policy-map route-nh-template <i>template-name</i> no lfa-policy-map
Context	config>service>vprn>isis>interface
Description	This command applies a route next-hop policy template to the IS-IS interface for the VPRN instance. When a route next-hop policy template is applied to an interface in IS-IS, it is applied in both level 1 and level 2. When a route next-hop policy template is applied to an interface in OSPF, it is applied in all areas. However, the command in an OSPF interface context can only be executed under the area in which the specified interface is primary and then applied in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command fails. If the user excluded the interface from LFA using the command loopfree-alternate-exclude, the LFA policy, if applied to the interface, has no effect. Finally, if the user applied a route next-hop policy template to a loopback interface or to the system interface, the command will not be rejected, but it will result in no action being taken.

The **no** form deletes the mapping of a route next-hop policy template to an OSPF or IS-IS interface.

Parameters *template-name* — Specifies the name of the template, up to 32 characters.

loopfree-alternate-exclude

Syntax **[no] loopfree-alternate-exclude**

Context config>service>vprn>isis>interface
 config>service>vprn>isis>level

Description This command instructs IGP to not include a specific interface or all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.

When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When it is excluded from the LFA SPF in OSPF, it is excluded in all areas. However, the above OSPF command can only be executed under the area in which the specified interface is primary and once enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command fails.

The **no** form of this command re-instates the default value for this command.

Default no loopfree-alternate-exclude

load-balancing-weight

Syntax **load-balancing-weight weight**
 no load-balancing-weight

Context config>service>vprn>isis>interface

Description This command configures the weighted ECMP load-balancing weight for an IS-IS interface of the VPRN. If the interface becomes an ECMP next-hop for IPv4 or IPv6 route and all the other ECMP next-hops are interfaces with configured (non-zero) load-balancing weights, then the traffic distribution over the ECMP interfaces is proportional to the weights. In other words, the interface with the largest load-balancing-weight should receive the most forwarded traffic if weighted ECMP is applicable.

The **no** form of this command disables weighted ECMP for the interface and, therefore, effectively disables weighted ECMP for any IP prefix that has this interface as a next-hop.

Default no load-balancing-weight

Parameters *weight* — Specifies the weight.

Values 0 to 4294967295

lsp-pacing-interval

Syntax	lsp-pacing-interval <i>milliseconds</i> no lsp-pacing-interval
Context	config>service>vprn>isis>interface
Description	This command configures the interval at which LSPs are sent from the interface. To avoid overwhelming neighbors that have less CPU processing power with LSPs, the pacing interval can be configured to limit how many LSPs are sent at the interval. LSPs are sent in bursts at the interval up to the configured limit. If a value of 0 is configured, no LSPs are sent from the interface. The no form of the command reverts to the default value. → Note: The IS-IS pacing interval is 100 milliseconds for values < 100 milliseconds, and 1 second for values ≥ 100 milliseconds. For example, a pacing interval of 2 milliseconds means that a maximum of 50 LSPs are sent in a burst at 100 millisecond intervals. The default pacing interval of 100 milliseconds means that a maximum of 10 LSPs are sent in a burst at 1 second intervals.
Default	lsp-pacing-interval 100 — the pacing interval is 100 milliseconds.
Parameters	<i>milliseconds</i> — Specifies the pacing interval in milliseconds at which IS-IS LSPs are sent from the interface at each interval expressed as a decimal integer. Values 0 to 65535

mesh-group

Syntax	mesh-group { <i>value</i> blocked } no mesh-group
Context	config>service>vprn>isis>interface
Description	This command assigns an interface to a mesh group. Mesh groups limit the amount of flooding that occurs when a new or changed LSP is advertised throughout an area. All routers in a mesh group should be fully meshed. When LSPs need to be flooded, only a single copy is received rather than a copy per neighbor. To create a mesh group, configure the same mesh group value for each interface that is part of the mesh group. All routers must have the same mesh group value configured for all interfaces that are part of the mesh group.

To prevent an interface from flooding LSPs, the optional **blocked** parameter can be specified. Configure mesh groups carefully. It is easy to create isolated islands that do not receive updates as (other) links fail.

The **no** form of the command removes the interface from the mesh group.

- Default** no mesh-group — The interface does not belong to a mesh group.
- Parameters** *value* — Specifies a unique decimal integer value distinguishes this mesh group from other mesh groups on this or any other router that is part of this mesh group.
- Values** 1 to 2000000000
- blocked** — Prevents an interface from flooding LSPs.

retransmit-interval

- Syntax** **retransmit-interval** *seconds*
no retransmit-interval
- Context** config>service>vprn>isis>interface
- Description** This command configures the minimum time between LSP PDU retransmissions on a point-to-point interface.
- The **no** form of the command reverts to the default value.
- Default** retransmit-interval 5
- Parameters** *seconds* — Specifies the interval in seconds that IS-IS LSPs can be sent on the interface
1 to 65535.

tag

- Syntax** **tag** *tag*
no tag
- Context** config>service>vprn>isis>interface
- Description** This command configures a route tag to the specified IP address of an interface.
- Parameters** *tag* — Specifies the tag value.
- Values** 1 to 4294967295

ipv4-multicast-routing

- Syntax** **ipv4-multicast-routing** {**native** | **mt**}

[no] ipv4-multicast-routing

Context	config>service>vprn>isis
Description	The multicast RTM is used for Reverse Path Forwarding checks. This command controls which IS-IS topology is used to populate the IPv4 multicast RTM. The no ipv4-multicast-routing form of the command results in none of the IS-IS routes being populated in the IPv4 multicast RTM and would be used if multicast is configured to use the unicast RTM for the RPF check.
Default	ipv4-multicast-routing native
Parameters	native — Causes IPv4 routes from the MT0 topology to be added to the multicast RTM for RPF checks. mt — Causes IPv4 routes from the MT3 topology to be added to the multicast RTM for RPF checks.

ipv4-routing

Syntax	[no] ipv4-routing
Context	config>service>vprn>isis
Description	This command specifies whether this IS-IS instance supports IPv4. The no form of the command disables IPv4 on the IS-IS instance.
Default	ipv4-routing

ipv6-routing

Syntax	[no] ipv6-routing {native mt}
Context	config>service>vprn>isis
Description	This command enables IPv6 routing. The no form of the command disables support for IS-IS IPv6 TLVs for IPv6 routing.
Default	no ipv6-routing
Parameters	native — Enables IS-IS IPv6 TLVs for IPv6 routing and enables support for native IPv6 TLVs. mt — Enables IS-IS multi-topology TLVs for IPv6 routing. When this parameter is specified, the support for native IPv6 TLVs is disabled.

level

Syntax	level <i>level-number</i>
Context	config>service>vprn>isis> config>service>vprn>isis>interface config>service>vprn>isis>link-group
Description	This command creates the context to configure IS-IS Level 1 or Level 2 area attributes. A router can be configured as a Level 1, Level 2, or Level 1-2 system. A Level 1 adjacency can be established if there is at least one area address shared by this router and a neighbor. A Level 2 adjacency cannot be established over this interface. Level 1/2 adjacency is created if the neighbor is also configured as Level 1/2 router and has at least one area address in common. A Level 2 adjacency is established if there are no common area IDs. A Level 2 adjacency is established if another router is configured as Level 2 or a Level 1/2 router with interfaces configured as Level 1/2 or Level 2. Level 1 adjacencies will not be established over this interface. To reset global and/or interface level parameters to the default, the following commands must be entered independently: <pre>level> no hello-authentication-key level> no hello-authentication-type level> no hello-interval level> no hello-multiplier level> no metric level> no passive level> no priority</pre>
Default	level 1 or level 2
Special Cases	Global IS-IS Level — The config>router>isis context configures default global parameters for both Level 1 and Level 2 interfaces. IS-IS Interface Level — The config>router>isis>interface context configures IS-IS operational characteristics of the interface at Level 1 and/or Level 2. A logical interface can be configured on one Level 1 and one Level 2. In this case, each level can be configured independently and parameters must be removed independently. By default an interface operates in both Level 1 and Level 2 modes.
Parameters	<i>level-number</i> — The IS-IS level number. Values 1, 2

default-ipv4-multicast-metric

Syntax	default-ipv4-multicast-metric <i>metric</i> no default-ipv4-multicast-metric
Context	config>service>vprn>isis>level
Description	This command configures the default metric to be used for the IS-IS interface in the IPv4 multicast topology (MT3). The no form of this command deletes the specified default metric and reverts to using the system default of 10.
Default	10
Parameters	<i>metric</i> — Specifies the default metric for interfaces in the IPv4 multicast topology (MT3). Values 1 to 16777215

default-ipv6-multicast-metric

Syntax	default-ipv6-multicast-metric <i>metric</i> no default-ipv6-multicast-metric
Context	config>service>vprn>isis>level
Description	This command configures the default metric to be used for the IS-IS interface in the IPv6 multicast topology (MT4). The no form of this command deletes the specified default metric and reverts to using the system default of 10.
Default	default-ipv6-multicast-metric 10
Parameters	<i>metric</i> — Specifies the default metric for interfaces in the IPv4 multicast topology (MT4). 1 to 16777215

default-ipv6-unicast-metric

Syntax	default-ipv6-unicast-metric <i>ipv6 metric</i> no default-ipv6-unicast-metric
Context	config>service>vprn>isis>level
Description	This command specifies the default metric for IPv6 unicast.
Default	default-ipv6-unicast-metric 10

Parameters	<i>ipv6-metric</i> — Specifies the default metric for IPv6 unicast.
Values	1 to 16777215

default-metric

Syntax	default-metric <i>ipv4 metric</i> no default-metric
Context	config>service>vprn>isis>level
Description	This command specifies the configurable default metric used for all IS-IS interfaces on this level. This value is not used if a metric is configured for an interface.
Default	default-metric 10
Parameters	<i>ipv4 metric</i> — Specifies the default metric for IPv4 unicast.
Values	1 to 16777214

external-preference

Syntax	external-preference <i>preference</i> no external-preference
Context	config>service>vprn>isis>level
Description	This command configures the external route preference for the IS-IS level. The external-preference command configures the preference level of either IS-IS level 1 or IS-IS level 2 external routes. By default, the preferences are as listed in the table below. A route can be learned by the router by different protocols, in which case, the costs are not comparable. When this occurs, the preference decides the route to use. Different protocols should not be configured with the same preference, if this occurs the tiebreaker is dependent on the default preference table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of the route to use is determined by the configuration of the ecmp in the config>router context.
Default	Default preferences are listed in Table 45 .

Table 45 **Default Preferences**

Route Type	Preference	Configurable
Direct attached	0	No
Static route	5	Yes
MPLS	7	—
OSPF internal routes	10	No
IS-IS Level 1 internal	15	Yes ¹
IS-IS Level 2 internal	18	Yes ¹
OSPF external	150	Yes
IS-IS Level 1 external	160	Yes
IS-IS Level 2 external	165	Yes
BGP	170	Yes
BGP	170	Yes

Note:

1. Internal preferences are changed using the **preference** command in the **config>router>isis>level level-number** context.

Parameters *preference* — The preference for external routes at this level as expressed.

Values 1 to 255

preference

Syntax **preference** *preference*
no preference

Context config>service>vprn>isis>level

Description This command configures the preference level of either IS-IS Level 1 or IS-IS Level 2 internal routes. By default, the preferences are listed in the table below.

A route can be learned by the router by different protocols, in which case, the costs are not comparable. When this occurs, the preference is used to decide to which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the table below. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision what route to use is determined by the configuration of the **ecmp** in the `config>router` context.

Default Default preferences are listed in [Table 46](#)

Table 46 Default Preferences

Route Type	Preference	Configurable
Direct attached	0	No
Static route	5	Yes
MPLS	7	—
OSPF internal routes	10	No
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes ¹
IS-IS level 2 external	165	Yes ¹
BGP	170	Yes

Note:

1. External preferences are changed using the **external-preference** command in the `config>router>isis>level level-number` context.

Parameters *preference* — The preference for external routes at this level expressed as a decimal integer.

Values 1 to 255

wide-metrics-only

Syntax `[no] wide-metrics-only`

Context `config>service>vprn>isis>level`

Description This command enables the exclusive use of wide metrics in the LSPs for the level number. Narrow metrics can have values between 1 and 63. IS-IS can generate two TLVs, one for the adjacency and one for the IP prefix. In order to support traffic engineering, wider metrics are required. When wide metrics are used, a second pair of TLVs are added, again, one for the adjacency and one for the IP prefix.

By default, both sets of TLVs are generated. When wide-metrics-only is configured, IS-IS only generates the pair of TLVs with wide metrics for that level.

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

level-capability

Syntax **level-capability** {**level-1** | **level-2** | **level-1/2**}
no level-capability

Context config>service>vprn>isis
config>service>vprn>isis>interface

Description This command configures the routing level for an instance of the IS-IS routing process.

An IS-IS router and an IS-IS interface can operate at Level 1, Level 2 or both Level 1 *and* 2.

[Table 47](#) displays configuration combinations and the potential adjacencies that can be formed.

Table 47 Potential Adjacency Capabilities

Global Level	Interface Level	Potential Adjacency
L 1/2	L 1/2	Level 1 and/or Level 2
L 1/2	L 1	Level 1 only
L 1/2	L 2	Level 2 only
L 2	L 1/2	Level 2 only
L 2	L 2	Level 2 only
L 2	L 1	none
L 1	L 1/2	Level 1 only
L 1	L 2	none
L 1	L 1	Level 1 only

The **no** form of the command removes the level capability from the configuration.

Default level-capability level-1/2

Special Cases	IS-IS Router — In the config>router>isis context, changing the level-capability performs a restart on the IS-IS protocol instance. IS-IS Interface — In the config>router>isis>interface context, changing the level-capability performs a restart of IS-IS on the interface.
Parameters	level-1 — Specifies the router/interface can operate at Level 1 only. level-2 — Specifies the router/interface can operate at Level 2 only. level-1/2 — Specifies the router/interface can operate at both Level 1 and Level 2.

link-group

Syntax	[no] link-group <i>link-group-name</i>
Context	config>service>vprn>isis
Description	This command configures a link-group for the router or VPRN instance. The no form of the command removes the specified link-group.
Parameters	<i>link-group-name</i> — Name of the link-group to be added or removed from the router or VPRN service.

description

Syntax	description <i>string</i> no description
Context	config>service>vprn>isis>link-group
Description	This command adds a description string to the associated link-group. The string can be up to 256 characters long and can only contain printable characters. If the command is issued in the context of a link-group that already contains a description then the previous description string is replaced. The no form of the command removes the description from the associated link-group.
Default	no description
Parameters	<i>string</i> — Character string to be associated with the associated link-group.

ipv4-multicast-metric-offset

Syntax	ipv4-multicast-metric-offset <i>offset-value</i> no ipv4-multicast-metric-offset

Context	config>service>vprn>isis>link-group>level
Description	This command sets the offset value for the IPv4 multicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric for the IPv4 multicast topology. The no form of the command reverts the offset value to 0.
Default	no ipv4-multicast-metric-offset
Parameters	<i>offset-value</i> — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold. Values 0 to 6777215

ipv4-unicast-metric-offset

Syntax	ipv4-unicast-metric-offset <i>offset-value</i> no ipv4-unicast-metric-offset
Context	config>service>vprn>isis>link-group>level
Description	This command sets the offset value for the IPv4 unicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric. The no form of the command reverts the offset value to 0.
Default	no ipv4-unicast-metric-offset
Parameters	<i>offset-value</i> — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold. Values 0 to 6777215

ipv6-unicast-metric-offset

Syntax	ipv6-unicast-metric-offset <i>offset-value</i> no ipv6-unicast-metric-offset
Context	config>service>vprn>isis>link-group>level
Description	This command sets the offset value for the IPv6 unicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric for the IPv6 topology. The no form of the command reverts the offset value to 0.

Default	no ipv6-unicast-metric-offset
Parameters	<i>offset-value</i> — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold.
Values	0 to 6777215

member

Syntax	[no] member <i>interface-name</i>
Context	config>service>vprn>isis>link-group>level
Description	This command adds or removes a links to the associated link-group. The interface name should already exist before it is added to a link-group. The no form of the command removes the specified interface from the associated link-group.
Parameters	<i>interface-name</i> — Name of the interface to be added or removed from the associated link-group.

oper-members

Syntax	oper-members <i>oper-members</i> no oper-members
Context	config>service>vprn>isis>link-group>level
Description	This command sets the threshold for the minimum number of operational links for the associated link-group. If the number of operational links drops below this threshold, the configured offsets are applied. For example, oper-members=3. The metric of the member interfaces is increased when the number of interfaces is lower than 3. The no form of the command reverts the oper-members limit to 1.
Default	no oper-members
Parameters	<i>oper-members</i> — Specifies the number of operational members.
Values	0 to 8

revert-members

Syntax	revert-members <i>revert-members</i> no revert-members
Context	config>router>isis>link-group

config>service>vprn>isis>link-group>level

Description This command sets the threshold for the minimum number of operational links to return the associated link-group to its normal operating state and remove the associated offsets to the IS-IS metrics. If the number of operational links is equal to or greater than the configured revert-member threshold then the configured offsets are removed.

The **no** form of the command reverts the revert-members threshold back to the default which is equal to the oper-member threshold value.

Parameters *revert-members* — Specifies the number of revert members.

Values 0 to 8

loopfree-alternates

Syntax [no] loopfree-alternates

Context config>service>vprn>isis

Description This command enables Loop-Free Alternate (LFA) computation by SPF under the IS-IS routing protocol level or under the OSPF routing protocol instance level.

When this command is enabled, it instructs the IGP SPF to attempt to pre-compute both a primary next-hop and an LFA next-hop for every learned prefix. When found, the LFA next-hop is populated into the routing table along with the primary next-hop for the prefix.

The **no** form of this command disables the LFA computation by IGP SPF.

Default no loopfree-alternates

exclude

Syntax exclude

Context config>service>vprn>isis>loopfree-alternates

Description This command excludes from LFA SPF calculation prefixes that match a prefix entry or a tag entry in a prefix policy.

The implementation already allows the user to exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.

If a prefix is excluded from LFA, then it will not be included in LFA calculation regardless of its priority. The prefix tag will, however, be used in the main SPF.



Note: Prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.

The default action of the **exclude** command, when not explicitly specified by the user in the prefix policy, is a “reject”. Thus, regardless if the user did or did not explicitly add the statement “default-action reject” to the prefix policy, a prefix that did not match any entry in the policy will be accepted into LFA SPF.

The **no** form of this command deletes the exclude prefix policy.

Default no exclude

prefix-policy

Syntax	prefix-policy <i>prefix-policy</i> [<i>prefix-policy</i>] no prefix-policy
Context	config>service>vprn>isis>loopfree-alternates>exclude
Description	This command specifies the name of the policy for the prefixes to exclude from the LFA SPF calculation in this ISIS instance. The no form of this command deletes the exclude prefix policy.
Default	no prefix-policy
Parameters	prefix-policy <i>prefix-policy</i> — Specifies the name of the prefix policy, up to 32 characters. Up to five prefix policies can be specified. The specified name must have been already defined.

lsp-lifetime

Syntax	lsp-lifetime <i>seconds</i> no lsp-lifetime
Context	config>service>vprn>isis
Description	This command sets the time, in seconds, the router wants the LSPs it originates to be considered valid by other routers in the domain. Each LSP received is maintained in an LSP database until the lsp-lifetime expires unless the originating router refreshes the LSP. By default, each router refreshes its LSPs every 20 minutes (1200 seconds) so other routers will not age out the LSP. The LSP refresh timer is derived from this formula: $\text{lsp-lifetime}/2$

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

Default	<code>lsp-lifetime 1200</code> — LSPs originated by the router should be valid for 1200 seconds (20 minutes).
Parameters	<i>seconds</i> — The time, in seconds, that the router wants the LSPs it originates to be considered valid by other routers in the domain.
Values	350 to 65535

lsp-mtu-size

Syntax	<code>lsp-mtu-size size</code> <code>no lsp-mtu-size</code>
Context	<code>config>service>vprn>isis</code> <code>config>service>vprn>isis>level</code>
Description	This command configures the LSP MTU size. If the <i>size</i> value is changed from the default using CLI or SNMP, then ISIS must be restarted for the change to take effect. This can be done by performing a shutdown command and then a no shutdown command in the config>router>isis context.



Note: Using the **exec** command to execute a configuration file to change the LSP MTU size from its default value will automatically restart IS-IS for the change to take effect.

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

Default	<code>lsp-mtu-size 1492</code>
Parameters	<i>size</i> — Specifies the LSP MTU size.
Values	490 to 9190

lsp-refresh-interval

Syntax	<code>lsp-refresh-interval [seconds] [half-lifetime {enable disable}]</code> <code>no lsp-refresh-interval</code>
Context	<code>config>service>vprn>isis</code>
Description	This command configures the IS-IS LSP refresh timer interval for the VPRN instance. When configuring the LSP refresh interval, the value that is specified for lsp-lifetime must also be considered. The LSP refresh interval cannot be greater than 90% of the LSP lifetime.

The **no** form of the command reverts to the default (600 seconds), unless this value is greater than 90% of the LSP lifetime. For example, if the LSP lifetime is 400, then the **no lsp-refresh-interval** command will be rejected.

Default	<code>lsp-refresh-interval 600 half-lifetime enable</code>
Parameters	<code>seconds</code> — Specifies the refresh interval. Values 150 to 65535 half-lifetime — Sets the refresh interval to always be half the lsp-lifetime value. When this parameter is set to enable , the configured refresh interval is ignored. Values enable, disable

multi-topology

Syntax	<code>[no] multi-topology</code>
Context	<code>config>service>vprn>isis</code>
Description	This command enables IS-IS multi-topology support.
Default	no multi-topology

ipv4-multicast

Syntax	<code>[no] ipv4-multicast</code>
Context	<code>config>service>vprn>isis>multi-topology</code>
Description	This command enables support for the IPv4 topology (MT3) within the associate IS-IS instance. The no form of this command disables support for the IPv4 topology (MT3) within the associated IS-IS instance.
Default	no ipv4-multicast

ipv6-unicast

Syntax	<code>[no] ipv6-unicast</code>
Context	<code>config>service>vprn>isis>multi-topology</code>
Description	This command enables multi-topology TLVs. The no form of the command disables multi-topology TLVs.

multicast-import

Syntax	[no] multicast-import
Context	config>service>vprn>isis
Description	This command enables ISIS to submit routes into the multicast Route Table Manager (RTM). The no form of the command disables the submission of routes into the multicast RTM.
Default	no multicast-import

overload

Syntax	overload [timeout seconds] [max-metric] no overload
Context	config>service>vprn>isis
Description	This command administratively sets the IS-IS router to operate in the overload state for a specific time period, in seconds, or indefinitely. During normal operation, the router may be forced to enter an overload state due to a lack of resources. When in the overload state, the router is only used if the destination is reachable by the router and will not be used for other transit traffic. If a time period is specified, the overload state persists for the configured length of time. If no time is specified, the overload state operation is maintained indefinitely. The overload command can be useful in circumstances where the router is overloaded or used prior to executing a shutdown command to divert traffic around the router. The max-metric parameter can be set to advertise transit links with the maximum metric of 0xfffffe (wide metrics) or 0x3f (regular metrics), instead of setting the overload bit when placing the router in overload. The no form of the command causes the router to exit the overload state.
Default	no overload
Parameters	seconds — The time, in seconds, that this router must operate in overload state. Values 60 to 1800 Default infinity (overload state maintained indefinitely) max-metric — Set the maximum metric instead of overload.

overload-export-external

Syntax	[no] overload-export-external
Context	config>service>vprn>isis
Description	This command enables external routes that are exported with an IS-IS export policy to continue to be advertised when the router is in overload. The no form of the command causes external routes to be withdrawn when the router is in overload.
Default	no overload-export-external

overload-export-interlevel

Syntax	[no] overload-export-interlevel
Context	config>service>vprn>isis
Description	This command enables inter-level routes that are exported with an IS-IS export policy to continue to be advertised when the router is in overload. The no form of the command causes inter-level routes to be withdrawn when the router is in overload.
Default	no overload-export-interlevel

overload-on-boot

Syntax	overload-on-boot [timeout seconds] [max-metric] no overload-on-boot
Context	config>service>vprn>isis
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: Step 1. The timeout timer expires. Step 2. A manual override of the current overload state is entered with the config>router>isis>no overload command. The no overload command does not affect the overload-on-boot function. If no timeout is specified, IS-IS will go into overload indefinitely after a reboot. After the reboot, the IS-IS status will display a permanent overload state: • L1 LSDB Overload : Manual on boot (Indefinitely in overload)

- L2 LSDB Overload : Manual on boot (Indefinitely in overload)

This state can be cleared with the **config>router>isis>no overload** command.

When specifying a timeout value, IS-IS will go into overload for the configured timeout after a reboot. After the reboot, the IS-IS status will display the remaining time the system stays in overload:

- L1 LSDB Overload : Manual on boot (Overload Time Left : 17)
- L2 LSDB Overload : Manual on boot (Overload Time Left : 17)

The overload state can be cleared before the timeout expires with the **config>router>isis>no overload** command.

The **no** form of the command removes the overload-on-boot functionality from the configuration.

Use the **show router isis status** command to display the administrative and operational state as well as all timers.

Default	no overload-on-boot
Parameters	timeout seconds — Configure the timeout timer for overload-on-boot in seconds.
	Values 60 to 1800
	max-metric — Set the maximum metric instead of overload.

poi-tlv-enable

Syntax	poi-tlv-enable no poi-tlv-enable
Context	config>service>vprn>isis
Description	Enable use of Purge Originator Identification (POI) TLV for this IS-IS instance. The POI is added to purges and contains the system ID of the router that generated the purge, which simplifies troubleshooting and determining what caused the purge. The no form of the command removes the POI functionality from the configuration.
Default	no poi-tlv-enable

prefix-attributes-tlv

Syntax	[no] prefix-attributes-tlv
Context	config>service>vprn>isis

Description	This command enables IS-IS Prefix Attributes TLV support to exchange extended IPv4 and IPv6 reachability information. Extended reachability information is required for traffic engineering features using path computation element (PCE) or optimal route reflection. The no form of the command removes the prefix-attributes-tlv configuration.
Default	no prefix-attributes-tlv

psnp-authentication

Syntax	[no] psnp-authentication
Context	config>service>vprn>isis config>service>vprn>isis>level
Description	This command enables authentication of individual ISIS packets of partial sequence number PDU (PSNP) type. The no form of the command suppresses authentication of PSNP packets.

prefix-limit

Syntax	prefix-limit <i>limit</i> [log-only] [threshold <i>percent</i>] [overload-timeout { <i>seconds</i> forever }] no prefix-limit
Context	config>service>vprn>isis
Description	This command configures the maximum number of prefixes that IS-IS can learn, and use to protect the system from a router that has accidentally advertised a large number of prefixes. If the number of prefixes reaches the configured percentage of this limit, an SNMP trap is sent. If the limit is exceeded, IS-IS will go into overload. The overload-timeout option controls the length of time that IS-IS is in the overload state when the prefix-limit is reached. The system automatically attempts to restart IS-IS at the end of this duration. If the overload-timeout forever option is used, IS-IS is not restarted automatically and stays in overload until the condition is manually cleared by the administrator. This is also the default behavior when the overload-timeout option is not configured. The no form of the command removes the prefix-limit .
Default	forever
Parameters	log-only — Enables a warning message to be sent at the specified threshold percentage and also when the limit is exceeded. However, overload is not set when this parameter is configured.

limit — Specifies the number of prefixes that can be learned, expressed as a decimal integer.

Values 1 to 4294967296

percent — Specifies the threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 to 100

seconds — Specifies the time in minutes before IS-IS is restarted.

Values 1 to 1800

forever — Specifies that IS-IS should be restarted only after the execution of the **clear router isis overload prefix-limit** command.

reference-bandwidth

Syntax	reference-bandwidth <i>bandwidth-in-kbps</i> reference-bandwidth [tbps <i>Tera-bps</i>] [gbps <i>Giga-bps</i>] [mbps <i>Mega-bps</i>] [kbps <i>Kilo-bps</i>] no reference-bandwidth
Context	config>service>vprn>isis
Description	This command configures the reference bandwidth that provides the basis of bandwidth relative costing. In order to calculate the lowest cost to reach a specific destination, each configured level on each interface must have a cost. If the reference bandwidth is defined, then the cost is calculated using the following formula: $\text{cost} = \text{reference-bandwidth} \div \text{bandwidth}$ If the reference bandwidth is configured as 10 Gigabits (10,000,000,000), a 100 M/bps interface has a default metric of 100. In order for metrics in excess of 63 to be configured, wide metrics must be deployed. (See wide-metrics-only in the config>router>isis context.) If the reference bandwidth is not configured, then all interfaces have a default metric of 10. The no form of the command reverts to the default value.
Default	no reference-bandwidth — No reference bandwidth is defined. All interfaces have a metric of 10.
Parameters	<i>bandwidth-in-kbps</i> — Specifies the reference bandwidth in kilobits per second expressed as a decimal integer. Values 1 to 4000000000 tbps <i>Tera-bps</i> — Specifies the reference bandwidth in terabits per second expressed as a decimal integer. Values 1 to 4

gbps *Giga-bps* — Specifies the reference bandwidth in Gigabits per second expressed as a decimal integer.

Values 1 to 999

mbps *Mega-bps* — Specifies the reference bandwidth in Megabits per second expressed as a decimal integer.

Values 1 to 999

kbps *Kilo-bps* — Specifies the reference bandwidth in kilobits per second expressed as a decimal integer.

Values 1 to 999

rib-priority

Syntax **rib-priority** {**high**} *prefix-list-name* | **tag** *tag-value*
no rib-priority

Context config>service>vprn>isis

Description This command enabled RIB prioritization for the IS-IS protocol and specifies the prefix list or IS-IS tag value that will be used to select the specific routes that should be processed through the IS-IS route calculation process at a higher priority.

The no rib-priority form of command disables RIB prioritization.

Default no rib-priority

Parameters *prefix-list-name* — Specifies the prefix list which is used to select the routes that are processed at a higher priority through the route calculation process.

tag tag-value — Specifies the tag value that is used to match IS-IS routes that are to be processed at a higher priority through the route calculation process.

Values 1 to 4294967295

router-id

Syntax **router-id** *ip-address*
no router-id

Context config>service>vprn>isis

Description This command sets the router ID for a specific VPRN context.

If neither the router ID nor system interface are defined, the router ID from the base router context is inherited.

The **no** form of the command removes the router ID definition from the given VPRN context.

Default	no router-id
Parameters	<i>ip-address</i> — The IP address must be given in dotted decimal notation.

standard-multi-instance

Syntax	[no] standard-multi-instance
Context	config>service>vprn>isis
Description	This command enables IS-IS multi-instance (MI) as described in draft-ginsberg-isis-mi-bis-01. Multiple instances allow instance-specific adjacencies to be formed that support multiple network topologies on the same physical interfaces. Each instance has an LSDB, and each PDU contains a TLV identifying the instance and the topology to which the PDU belongs. A single topology is supported in each instance, so the instance-specific topology identifier (ITID) is set to 0 and cannot be changed. The standard-multi-instance (based on draft-ginsberg-isis-mi-bis-01) and iid-tlv-enable (based on draft-ietf-isis-mi-02) commands cannot be configured in the same instance, because the MAC addresses and PDUs from the two standards are incompatible. The no form of the command removes the standard-multi-instance configuration.
Default	no standard-multi-instance

timers

Syntax	[no] timers
Context	config>service>vprn>isis
Description	This command configures the IS-IS timer values.
Default	n/a

lsp-wait

Syntax	lsp-wait <i>lsp-wait</i> [lsp-initial-wait <i>initial-wait</i>] [lsp-second-wait <i>second-wait</i>]
Context	config>service>vprn>isis>timers
Description	This command is used to customize LSP generation throttling. Timers that determine when to generate the first, second, and subsequent LSPs can be controlled with this command. Subsequent LSPs are generated at increasing intervals of the second lsp-wait timer until a maximum value is reached.



Note: The timer granularity is 10 ms if the value is < 500 ms, and 100 ms if the value is ≥ 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Parameters	<i>lsp-wait</i> — Specifies the maximum interval, in milliseconds, between two consecutive occurrences of an LSP being generated.
	Values 10 to 120000
	Default 5000
	<i>initial-wait</i> — Specifies the initial LSP generation delay, in milliseconds. Values less than 100 ms are internally rounded down to 0, so that there is no added initial LSP generation delay.
	Values 10 to 100000
	Default 10
	<i>second-wait</i> — Specifies the hold time, in milliseconds, between the first and second LSP generation.
	Values 10 to 100000
	Default 1000

spf-wait

Syntax	[no] spf-wait <i>spf-wait</i> [<i>spf-initial-wait initial-wait</i>] [<i>spf-second-wait second-wait</i>]
Context	config>service>vprn>isis>timers
Description	This command defines the maximum interval, in milliseconds, 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, then the next SPF will run after 2000 milliseconds, and the next SPF after that will run after 4000 milliseconds, and so on, until it reaches the <i>spf-wait</i> value. The SPF interval will stay at the <i>spf-wait</i> 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 the SPF <i>initial-wait</i> value.



Note: The timer granularity is 100 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Parameters	<i>spf-wait</i> — Specifies the maximum interval, in milliseconds, between two consecutive SPF calculations. Values 10 to 120000 Default 10000 <i>initial-wait</i> — Specifies the initial SPF calculation delay, in milliseconds, after a topology change. Values 10 to 100000 Default 1000 <i>second-wait</i> — Specifies the hold time, in milliseconds, between the first and second SPF calculation. Values 10 to 100000 Default 1000

strict-adjacency-check

Syntax	[no] strict-adjacency-check
Context	config>service>vprn>isis
Description	This command enables strict checking of address families (IPv4 and IPv6) for IS-IS adjacencies. When enabled, adjacencies will not come up unless both routers have exactly the same address families configured. If there is an existing adjacency with unmatched address families, it will be torn down. This command is used to prevent black-holing traffic when IPv4 and IPv6 topologies are different. When disabled (no strict-adjacency-check) a BFD session failure for either IPv4 or Ipv6 will cause the routes for the other address family to be removed as well. When disabled (no strict-adjacency-check), both routers only need to have one common address family to establish the adjacency.
Default	no strict-adjacency-check

summary-address

Syntax	summary-address {<i>ip-prefix/mask</i> <i>ip-prefix [netmask]</i>} [<i>level</i>] [tag <i>tag</i>] no summary-address {<i>ip-prefix/mask</i> <i>ip-prefix [netmask]</i>}
Context	config>service>vprn>isis
Description	This command creates summary-addresses for the specified router or VPRN instance.

Parameters *ip-prefix/mask* — Specifies information for the specified IP prefix and mask length.

Values

ip-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 to 32
ipv6-prefix	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D
ipv6-prefix-length	0 to 128

netmask — The subnet mask in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

level — Specifies IS-IS level area attributes. If no level parameter is specified, the default is level-1/2.

Values level-1, level-2, level-1/2

tag *tag* — Assigns a route tag to the summary address.

Values 1 to 4294967295

system-id

Syntax **system-id** *isis-system-id*

no system-id

Context config>service>vprn>isis

Description This command configures the IS-IS system ID. The system ID has a fixed length of 6 octets; it is determined using the following preference:

1. **config>service>vprn>isis>system-id**
2. **config>service>vprn>isis>router-id**
3. **config>service>vprn>router-id**
4. **config>service>vprn>if>address**
5. The default system ID 2550.0000.0000, based on the default router ID 255.0.0.0

The system ID is integral to IS-IS; therefore, for the **system-id** command to take effect, a **shutdown** and then **no shutdown** must be performed on the IS-IS instance. This will ensure that the configured and operational system ID are always the same.

The **no** form of the command removes the system ID from the configuration. The router ID is used when no system ID is specified.

Default no system-id

Parameters	<i>isis-system-id</i> — 12 hexadecimal characters in dotted-quad notation.
Values	aaaa.bbbb.cccc, where aaaa, bbbb, and cccc are hexadecimal numbers

ignore-attached-bit

Syntax	ignore-attached-bit no ignore-attached-bit
Context	config>service>vprn>isis
Description	This command configures IS-IS to ignore the attached bit on received Level 1 LSPs to disable installation of default routes.

suppress-attached-bit

Syntax	[no] suppress-attached-bit
Context	config>service>vprn>isis
Description	This command configures IS-IS to suppress setting the attached bit on originated Level 1 LSPs to prevent all L1 routers in the area from installing a default route to it.

import

Syntax	import <i>policy-name</i> [<i>policy-name</i> ... (up to 5 max)] no import
Context	config>service>vprn>isis
Description	This command applies one or more (up to five) route policies as IS-IS import policies. When a prefix received in an IS-IS LSP is accepted by an entry in an IS-IS import policy, it is installed in the routing table, if it is the most preferred route to the destination. When a prefix received in an IS-IS LSP is rejected by an entry in an IS-IS import policy, it is not installed in the routing table, even if it has the lowest preference value among all the routes to that destination. The flooding of LSPs is unaffected by IS-IS import policy actions. The no form of the command removes all policies from the configuration.
Default	no import

Parameters *policy-name* — Identify the export route policy name. Allowed values are any string up to 32 characters long, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. The specified name(s) must already be defined.

unicast-import-disable

Syntax **[no] unicast-import-disable**

Context config>service>vprn>isis

Description This command allows one IGP to import its routes into RPF RTM while another IGP imports routes only into the unicast RTM. Import policies can redistribute routes from an IGP protocol into the RPF RTM (the multicast routing table). By default, the IGP routes will not be imported into RPF RTM as such an import policy must be explicitly configured.

Default no unicast-import-disable

3.8.2.31 OSPF Commands

ospf

Syntax **ospf [router-id]**
 no ospf

Context config>service>vprn

Description This command enables access to the context to enable an OSPF protocol instance.

 OSPF instances are **shutdown** when created, so that all parameters can be configured prior to the instance being enabled.

 The **no** form of the command deletes the OSPF protocol instance removing all associated configuration parameters.

Default no ospf

Parameters *router-id* — Specifies the OSPF router ID to be used with the associated OSPF instance. The *router-id* must be given a dot decimal notation format.

Values a.b.c.d

ospf3

Syntax	ospf3 [<i>instance-id</i>] [<i>router-id</i>] [no] ospf3 <i>instance-id</i>
Context	config>service>vprn
Description	This command creates an OSPFv3 routing instance and then enters the associated context to configure associated protocol parameters. OSPF instances are shutdown when created, so that all parameters can be configured prior to the instance being enabled. The no form of the command deletes the OSPFv3 protocol instance, removing all associated configuration parameters.
Default	no ospf3
Parameters	<i>instance-id</i> — Specifies the instance ID for the OSPFv3 instance being created or modified. The instance ID must match the specified range based on the address family. For ipv6-unicast, the instance id must be between 0 and 31. For ipv4-unicast the instance id must be between 64 and 95. Values 0 to 31: IPV6 unicast 64 to 95: IPV4 unicast <i>router-id</i> — Specifies the IP address.

advertise-router-capability

Syntax	advertise-router-capability advertise-router-capability {link area as} no advertise-router-capability
Context	config>service>vprn>ospf config>service>vprn>ospf3 config>service>vprn>ospf>area config>service>vprn>ospf>area>interface config>service>vprn>ospf3>area>interface
Description	This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A Router Information (RI) LSA as defined in RFC 4970 advertises the following capabilities: • OSPF graceful restart capable: no • OSPF graceful restart helper: yes, when enabled • OSPF Stub Router support: yes • OSPF Traffic Engineering support: yes, when enabled • OSPF point-to-point over LAN: yes

- OSPF Experimental TE: no

The parameters (link, area & as) control the scope of the capabilities advertisements.

The no form of this command, disables this capability.

Default	no advertise-router-capability
Parameters	link — Capabilities are only advertised over local link and not flooded beyond. area — Capabilities are only advertised within the area of origin. as — Capabilities are only advertised throughout the entire autonomous system.

area

Syntax	[no] area <i>area-id</i>
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	This command creates the context to configure an OSPF area. An area is a collection of network segments within an AS that have been administratively grouped together. The area ID can be specified in dotted decimal notation or as a 32-bit decimal integer. The no form of the command deletes the specified area from the configuration. Deleting the area also removes the OSPF configuration of all the interfaces, virtual-links, sham-links, and address-ranges and so on, that are currently assigned to this area.
Default	no area — No OSPF areas are defined.
Parameters	<i>area-id</i> — 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 <i>ip-prefix/prefix-length</i> [advertise not-advertise] no area-range <i>ip-prefix/mask</i> area-range <i>ipv6-prefix/prefix-length</i> [advertise not-advertise] no area-range <i>ipv6-prefix/prefix-length</i>
Context	config>service>vprn>ospf>area config>service>vprn>ospf3>area config>service>vprn>ospf>area>nssa config>service>vprn>ospf3>area>nssa

Description This command creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression. When a range is created, the range is configured to be advertised or not advertised into other areas. Multiple range commands may be used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.

ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.

The **no** form of the command deletes the range (non) advertisement.

Default no area-range

Special Cases **NSSA Context** — In the NSSA context, the option specifies that the range applies to external routes (via type-7 LSAs) learned within the NSSA when the routes are advertised to other areas as type-5 LSAs.

Area Context — If this command is not entered under the NSSA context, the range applies to summary LSAs even if the area is an NSSA.

Parameters *ipv6-prefix/prefix-length* — The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.

Values

ipv6-prefix	x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
ipv6-prefix-length	0 to 128

mask — The subnet mask for the range expressed as a decimal integer mask length or in dotted decimal notation.

Values 0 to 32 (mask length), 0.0.0.0 to 255.255.255.255 (dotted decimal)

advertise | not-advertise — Specifies whether or not to advertise the summarized range of addresses into other areas. The **advertise** keyword indicates the range will be advertised, and the keyword **not-advertise** indicates the range will not be advertised.

The default is **advertise**.

blackhole-aggregate

Syntax [no] blackhole-aggregate

Context config>service>vprn>ospf>area
config>service>vprn>ospf3>area

Description This command installs a low priority blackhole route for the entire aggregate. Existing routes that make up the aggregate will have a higher priority and only the components of the range for which no route exists are blackholed.

It is possible that when performing area aggregation, addresses may be included in the range for which no actual route exists. This can cause routing loops. To avoid this problem configure the blackhole aggregate option.

The **no** form of this command removes this option.

Default blackhole-aggregate

export

Syntax **export** *policy-name* [*policy-name...*(up to 5 max)]
no export

Context config>service>vprn>ospf>area
config>service>vprn>ospf3>area

Description This command configures ABR export policies to filter OSPFv2 Type 3 Summary-LSAs or OSPFv3 Inter-Area-Prefix-LSA between areas, in order to only permit the specified routes from being exported into an area.

This command cannot be used in OSPF area 0.

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

Default no export

Parameters *policy-name* — The export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

The specified names must already be defined.

import

Syntax **import** *policy-name* [*policy-name...*(up to 5 max)]
no import

Context config>service>vprn>ospf>area
config>service>vprn>ospf3>area

Description This command configures ABR import policies to filter OSPFv2 Type 3 Summary-LSAs or OSPFv3 Inter-Area-Prefix-LSA between areas, in order to only permit the specified routes from being imported into an area.

This command cannot be used in OSPF area 0.

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

Default	no import
Parameters	<i>policy-name</i> — The export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. The specified names must already be defined.

interface

Syntax	interface <i>ip-int-name</i> [<i>secondary</i>] no interface <i>ip-int-name</i>
Context	config>service>vprn>ospf>area config>service>vprn>ospf3>area
Description	This command creates a context to configure an OSPF interface. By default interfaces are not activated in any interior gateway protocol such as OSPF unless explicitly configured. The no form of the command deletes the OSPF interface configuration for this interface. The shutdown command 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	<i>ip-int-name</i> — The IP interface name. Interface names must be unique within the group of defined IP interfaces for config router interface and config service vprn interface commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. If the IP interface name does not exist or does not have an IP address configured an error message will be returned. If the IP interface exists in a different area it will be moved to this area. <i>secondary</i> — Allows multiple secondary adjacencies to be established over a single IP interface.

sham-link

Syntax	sham-link <i>ip-int-name</i> <i>ip-address</i>

Context	config>service>vprn>ospf>area
Description	This command is similar to a virtual link with the exception that metric must be included in order to distinguish the cost between the MPLS-VP RN link and the backdoor.
Parameters	<i>ip-int-name</i> — The local interface name used for the sham-link. This is a mandatory parameter and interface names must be unique within the group of defined IP interfaces for config>router>interface, config>service>ies>interface and config>service>vprn>interface commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters, the entire string must be enclosed between double quotes. If the IP interface name does not exist or does not have an IP address configured, an error message will be returned. <i>ip-address</i> — The IP address of the sham-link neighbor in IP address dotted decimal notation. This parameter is the remote peer of the sham link's IP address used to set up the sham-link. This is a mandatory parameter and must be a valid IP address.

advertise-subnet

Syntax	[no] advertise-subnet
Context	config>service>vprn>ospf>area>if
Description	This command enables advertising point-to-point interfaces as subnet routes (network number and mask). When disabled, point-to-point interfaces are advertised as host routes. This command is not supported in the OSPF3 context. The no form of the command disables advertising point-to-point interfaces as subnet routes meaning they are advertised as host routes.
Default	advertise-subnet — Advertises point-to-point interfaces as subnet routes.

auth-keychain

Syntax	auth-keychain <i>name</i>
Context	config>service>vprn>ospf>area>if config>service>vprn>ospf>area>sham-link config>service>vprn>ospf>area>virtual-link
Description	This command enables the authentication keychain.
Parameters	<i>name</i> — Specifies the name of the authentication keychain, up to 32 characters.

authentication

Syntax	authentication bidirectional <i>sa-name</i> authentication inbound <i>sa-name</i> outbound <i>sa-name</i> no authentication
Context	config>service>vprn>ospf3>area>if config>service>vprn>ospf3>area>virtual-link
Description	This command configures OPSFv3 confidentiality authentication. The no form of the command removes the SA name from the configuration.
Parameters	bidirectional <i>sa-name</i> — Specifies the IPSec security association name in case the OSPFv3 traffic on the interface has to be authenticated. inbound <i>sa-name</i> — Specifies the IPSec security association name in case the OSPFv3 traffic on the interface has to be authenticated. outbound <i>sa-name</i> — Specifies the IPSec security association name in case the OSPFv3 traffic on the interface has to be authenticated.

authentication-key

Syntax	authentication-key { <i>authentication-key</i> <i>hash-key</i> } [hash hash2] no authentication-key
Context	config>service>vprn>ospf>area>if config>service>vprn>ospf>area>virtual-link config>service>vprn>ospf>area>sham-link
Description	This command configures the password used by the OSPF interface or virtual-link to send and receive OSPF protocol packets on the interface when simple password authentication is configured. This command is not valid in the OSPF3 context. All neighboring routers must use the same type of authentication and password for proper protocol communication. If the authentication-type is configured as password, then this key must be configured. By default, no authentication key is configured. This command is not supported in the OSPF context. The no form of the command removes the authentication key.
Default	no authentication-key — No authentication key is defined.

- Parameters**
- authentication-key* — The authentication key. The key can be any combination of ASCII characters up to 8 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").
 - hash-key* — The hash key. The key can be any combination of ASCII characters up to 22 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").
This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.
 - hash** — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.
 - hash2** — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

authentication-type

- Syntax** **authentication-type {password | message-digest}**
no authentication-type
- Context** config>service>vprn>ospf>area>if
 config>service>vprn>ospf>area>sham-link
 config>service>vprn>ospf>area>virtual-link
- Description** This command enables authentication and specifies the type of authentication to be used on the OSPF interface, virtual-link, and sham-link.

This command is not valid in the OSPF3 context.

Both simple **password** and **message-digest** authentication are supported.

The **no** form of the command disables authentication on the interface.
- Default** no authentication-type — No authentication is enabled on an interface.
- Parameters** **password** — This keyword enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple **password** authentication is enabled.

message-digest — This keyword enables message digest MD5 authentication in accordance with RFC1321. If this option is configured, then at least one message-digest-key must be configured.

bfd-enable

Syntax	bfd-enable [remain-down-on-failure] no bfd-enable
Context	config>service>vprn>ospf>interface config>service>vprn>ospf3>area>if
Description	This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface. The no form of this command removes BFD from the associated IGP protocol adjacency.
Default	no bfd-enable
Parameters	remain-down-on-failure — Forces adjacency down on BFD failure.

dead-interval

Syntax	dead-interval <i>seconds</i> no dead-interval
Context	config>service>vprn>ospf>area>if config>service>vprn>ospf>area>virtual-link config>service>vprn>ospf3>area>if config>service>vprn>ospf3>area>virtual-link config>service>vprn>ospf>area>sham-link
Description	This command configures the time, in seconds, that OSPF waits before declaring a neighbor router down. If no hello packets are received from a neighbor for the duration of the dead interval, the router is assumed to be down. The minimum interval must be two times the hello interval. The no form of the command reverts to the default value.
Default	dead-interval 40
Special Cases	OSPF Interface — If the dead-interval configured applies to an interface, then all nodes on the subnet must have the same dead interval. Virtual Link — If the dead-interval configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same dead interval. Sham-link — If the dead-interval configured applies to a sham-link, then the interval on both endpoints of the sham-link must have the same dead interval.

Parameters	<i>seconds</i> — The dead interval expressed as a decimal integer.
Values	2 to 2147483647 seconds

graceful-restart

Syntax	[no] graceful-restart
Context	config>service>vprn>ospf
Description	This command enables OSPF graceful restart (GR) to minimize service interruption. When the control plane of a GR-capable router fails or restarts, the neighboring routers (GR helpers) temporarily preserve OSPF forwarding information. Traffic continues to be forwarded to the restarting router using the last known forwarding tables. If the control plane of the restarting router becomes operationally and administratively up within the grace period, the restarting router resumes normal OSPF operation. If the grace period expires, then the restarting router is presumed inactive and the OSPF topology is recalculated to route traffic around the failure. The no form of the command disables graceful restart and removes the graceful restart configuration from the OSPF instance.
Default	no graceful-restart

helper-disable

Syntax	[no] helper-disable
Context	config>service>vprn>ospf>graceful-restart config>service>vprn>ospf3>graceful-restart
Description	This command disables helper support for OSPF graceful restart (GR). When graceful-restart is enabled, the router can be a helper (meaning that the router is helping a neighbor to restart), a restarting router, or both. The router only supports helper mode. It will not act as a restarting router, because the high availability feature set already preserves OSPF forwarding information such that this functionality is not needed. This command is a historical command and should not be disabled. Configuring helper-disable has the effect of disabling graceful restart, because the router only supports helper mode. The no helper-disable command enables helper support and is the default when graceful restart is enabled.
Default	no helper-disable

strict-lsa-checking

Syntax	[no] strict-lsa-checking
Context	config>service>vprn>ospf>graceful-restart config>service>vprn>ospf3>graceful-restart
Description	This command indicates whether an OSPF restart helper should terminate graceful restart when there is a change to an LSA that would be flooded to the restarting router during the restart process. The default OSPF behavior is to terminate a graceful restart if an LSA changes, which causes the OSPF neighbor to go down. The no strict-lsa-checking command disables strict LSA checking.
Default	strict-lsa-checking

ignore-dn-bit

Syntax	[no] ignore-dn-bit
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	This command specifies whether to suppress the setting of the DN bit for OSPF or OSPF3 LSA packets generated by this instance of OSPF or OSPF3 on the router. The no form of the command enables the OSPF or OSPF3 router to follow the normal procedure to determine whether to set the DN bit.
Default	no ignore-dn-bit

import

Syntax	import <i>policy-name</i> [<i>policy-name...</i> (up to 5 max)] no import
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	This command applies one or more (up to 5) route policies as OSPF import policies. When a prefix received in an OSPF LSA is accepted by an entry in an OSPF import policy it is installed in the routing table if it is the most preferred route to the destination. When a prefix received in an OSPF LSA is rejected by an entry in an OSPF import policy it is not installed in the routing table, even if it has the lowest preference value among all the routes to that destination. The flooding of LSAs is unaffected by OSPF import policy actions. This command only applies to the 7750 SR.

Default	If an OSPF route has the lowest preference value among all routes to a destination it is installed in the routing table.
Parameters	<i>policy-name</i> — The import route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. The specified name(s) must already be defined.

hello-interval

Syntax	hello-interval <i>seconds</i> no hello-interval
Context	<pre>config>service>vprn>ospf>area>if config>service>vprn>ospf3>area>if config>service>vprn>ospf>area>virtual-link config>service>vprn>ospf3>area>virtual-link config>service>vprn>ospf>area>sham-link</pre>
Description	This command configures the interval between OSPF hellos issued on the interface, virtual link, or sham-link. The hello interval, in combination with the dead-interval, is used to establish and maintain the adjacency. Use this parameter to edit the frequency that hello packets are sent. Reducing the interval, in combination with an appropriate reduction in the associated dead-interval, allows for faster detection of link and/or router failures at the cost of higher processing costs. The no form of this command reverts to the default value.
Default	hello-interval 10 — a 10-second hello interval
Special Cases	OSPF Interface — If the hello-interval configured applies to an interface, then all nodes on the subnet must have the same hello interval. Virtual Link — If the hello-interval configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same hello interval. Sham Link — If the hello-interval configured applies to a sham-link, then the interval on both endpoints of the sham-link must have the same hello interval.
Parameters	<i>seconds</i> — The hello interval in seconds expressed as a decimal integer. Values 1 to 65535

interface-type

Syntax	interface-type {broadcast point-to-point non-broadcast} no interface-type
Context	config>service>vprn>ospf>area>if config>service>vprn>ospf3>area>if
Description	This command configures the interface type to be one of broadcast, point-to-point, or non-broadcast. Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the broadcast adjacency maintenance overhead if the Ethernet link provided the link is used as a point-to-point. If the interface type is not known at the time the interface is added to OSPF and subsequently the IP interface is bound (or moved) to a different interface type, this command must be entered manually. The no form of the command reverts to the default value.
Default	point-to-point — If the physical interface is SONET. broadcast — If the physical interface is Ethernet or unknown.
Special Cases	Virtual-Link — A virtual link is always regarded as a point-to-point interface and is not configurable.
Parameters	broadcast — Configures the interface as a broadcast network. To significantly improve adjacency forming and network convergence, configure the network as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet. point-to-point — Configures the interface as a point-to-point link. non-broadcast — Configures the interface as a non-broadcast network.

loopfree-alternate-exclude

Syntax	[no] loopfree-alternate-exclude
Context	config>service>vprn>ospf>area config>service>vprn>ospf3>area config>service>vprn>ospf>area>interface config>service>vprn>ospf3>area>interface
Description	This command instructs IGP to not include a specific interface or all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.

When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When it is excluded from the LFA SPF in OSPF, it is excluded in all areas. However, the above OSPF command can only be executed under the area in which the specified interface is primary and once enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command fails.

The **no** form of this command re-instates the default value for this command.

Default no loopfree-alternate-exclude

lsa-filter-out

Syntax **lsa-filter-out** [**all** | **except-own-rtrlsa** | **except-own-rtrlsa-and-defaults**]
no lsa-filter-out

Context config>router>ospf>area>interface
config>router>ospf3>area>interface
config>service>vprn>ospf>area>interface
config>service>vprn>ospf3>area>interface

Description This command enables filtering of outgoing OSPF LSAs on the selected OSPFv2 or OSPFv3 interface. Three filtering options are provided:

- Do not flood any LSAs out the interface. This option is suitable if the neighbor is simply-connected and has a statically configured default route with the address of this interface as next-hop.
- Flood the router's own router-LSA out the interface and suppress all other flooded LSAs. This option is suitable if the neighbor is simply-connected and has a statically configured default route with a loopback or system interface address (contained in the router-LSA) as next-hop.
- Flood the router's own router-LSA and all self-generated type-3, type-5 and type-7 LSAs advertising a default route (0/0) out the interface; suppress all other flooded LSAs. This option is suitable if the neighbor is simply-connected and does not have a statically configured default route.

The **no** form of this command disables OSPF LSA filtering (normal operation).

Default no lsa-filter-out

multicast-import

Syntax [**no**] **multicast-import**

Context config>service>vprn>ospf
config>service>vprn>ospf3

Description	This command enables the submission of routes into the multicast Route Table Manager (RTM) by OSPF. The no form of the command disables the submission of routes into the multicast RTM.
Default	no multicast-import

message-digest-key

Syntax	message-digest-key <i>keyid</i> md5 { <i>key</i> <i>hash-key</i> } [<i>hash</i>] no message-digest-key <i>keyid</i>
Context	config>service>vprn>ospf>area>if config>service>vprn>ospf>area>virtual-link config>service>vprn>ospf>area>sham-link
Description	This command configures a message digest key when MD5 authentication is enabled on the interface, virtual-link or sham-link. Multiple message digest keys can be configured. This command is not valid in the OSPF3 context. The no form of the command removes the message digest key identified by the <i>key-id</i> .
Default	No message digest keys are defined.
Parameters	keyid — The <i>keyid</i> is expressed as a decimal integer. Values 1 to 255 md5 <i>key</i> — The MD5 key. The <i>key</i> can be any alphanumeric string up to 16 characters in length. md5 <i>hash-key</i> — The MD5 hash key. The key can be any combination of ASCII characters up to 32 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" "). This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided. hash — Specifies the key is entered in an encrypted form. If the hash parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash parameter specified.

metric

Syntax	metric <i>metric</i> no metric
Context	config>service>vprn>ospf>area>if

```
config>service>vprn>ospf3>area>if
config>service>vprn>ospf>area>sham-link
```

Description This command configures an explicit route cost metric for the OSPF interface that overrides the metrics calculated based on the speed of the underlying link.

The **no** form of the command deletes the manually configured interface metric, so the interface uses the computed metric based on the **reference-bandwidth** command setting and the speed of the underlying link.

Default no metric

Parameters *metric* — The metric to be applied to the interface expressed as a decimal integer.

Values 1 to 65535

mtu

Syntax *mtu bytes*
no mtu

Context config>service>vprn>ospf>area>if
config>service>vprn>ospf3>area>if

Description This command configures the OSPF packet size used on this interface. If this parameter is not configured OSPF derives the MTU value from the MTU configured (default or explicitly) in the following contexts:

```
config>port>ethernet
config>port>sonet-sdh>path
config>port>tdm>t3-e3
config>port>tdm>t1-e1>channel-group
```

If this parameter is configured, the smaller value between the value configured here and the MTU configured (default or explicitly) in an above-mentioned context is used.

To determine the actual packet size add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF (IP) packet MTU configured in this command.

Use the **no** form of this command to revert to default value derived from the MTU configured in the **config>port** context.

Default no mtu

Parameters *bytes* — 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 config>service>vprn>ospf3>area>if
Description	This command adds the passive property to the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol. By default, only interface addresses that are configured for OSPF will be advertised as OSPF interfaces. The passive parameter allows an interface to be advertised as an OSPF interface without running the OSPF protocol. While in passive mode, the interface will ignore ingress OSPF protocol packets and not transmit any OSPF protocol packets. The no form of the command removes the passive property from the OSPF interface.
Default	passive (service interfaces defined in config>router>service-prefix) no passive (all other interfaces)

neighbor

Syntax	[no] neighbor ip-address
Context	config>service>vprn>ospf>area>interface config>service>vprn>ospf3>area>interface
Description	This command configures an OSPF non-broadcast multi-access (NBMA) neighbor. The OSPF interface must be configured as an NBMA interface with the interface-type non-broadcast command. An NBMA network has no broadcast or multicast capabilities, so the router cannot discover its neighbors dynamically. All neighbors must be configured statically with the neighbor command. In addition to configuring the OSPF NBMA neighbor's IP address, the neighbor's MAC address may need to be configured with the config>service>vprn>interface>static-arp command for OSPFv2 neighbors using its IPv4 address, and the config>service>vprn>interface>ipv6>neighbor command for OSPFv3 neighbors using its IPv6 link-local address. The no form of the command removes the neighbor configuration.
Default	No OSPF NBMA neighbors are configured.

Parameters *ip-address* — Specifies the OSPFv2 neighbor's IPv4 address or the OSPFv3 neighbor's IPv6 link-local address.

Values

ipv4-address: a.b.c.d
 ipv6-address: x:x:x:x:x:x [-interface]
 x:x:x:x:x:d.d.d [-interface]
 x: [0..FFFF]H
 d: [0..255]D
 interface —32 characters max, for link local addresses.

poll-interval

Syntax **poll-interval** *seconds*
 no poll-interval

Context config>service>vprn>ospf>area>interface
 config>service>vprn>ospf3>area>interface

Description This command configures the poll interval, in seconds. The poll interval is the time between two Hello packets to a dead (non-adjacent) OSPF NBMA neighbor. The default value of the poll interval timer is higher than the hello interval timer to avoid wasting bandwidth on non-broadcast networks, since OSPF messages are unicast to each configured neighbor. The poll interval timer is used only on **non-broadcast** interface types and has no effect if configured on other interface types.

 The **no** form of the command removes the **poll-interval** configuration.

Default 120

Parameters *seconds* — Specifies the poll interval, in seconds.

Values 0 to 4294967295

priority

Syntax **priority** *number*
 no priority

Context config>service>vprn>ospf>area>if
 config>service>vprn>ospf3>area>if

Description This command configures the priority of the OSPF interface that is used to elect the designated router on the subnet.

This parameter is only used if the interface is of type **broadcast**. The router with the highest priority interface becomes the designated router. A router with priority 0 is not eligible to be the designated router or backup designated router.

The **no** form of the command resets the interface priority to the default value.

Default	priority 1
Parameters	<i>number</i> — The interface priority expressed as a decimal integer. A value of 0 indicates the router is not eligible to be the Designated Router or Backup Designated Router on the interface subnet.
Values	0 to 255

retransmit-interval

Syntax	retransmit-interval <i>seconds</i> no retransmit-interval
Context	config>service>vprn>ospf>area>if config>service>vprn>ospf>area>virtual-link config>service>vprn>ospf3>area>if config>service>vprn>ospf3>area>virtual-link config>service>vprn>ospf>area>sham-link
Description	This command specifies the length of time, in seconds, that OSPF will wait before retransmitting an unacknowledged link state advertisement (LSA) to an OSPF neighbor. The value should be longer than the expected round trip delay between any two routers on the attached network. Once the retransmit-interval expires and no acknowledgment has been received, the LSA will be retransmitted. The no form of this command reverts to the default interval.
Default	retransmit-interval 5
Parameters	<i>seconds</i> — The retransmit interval in seconds expressed as a decimal integer.
Values	1 to 3600

rib-priority

Syntax	rib-priority { <i>high</i> } <i>prefix-list-name</i> no rib-priority
Context	config>service>vprn>ospf>area>interface config>service>vprn>ospf3>area>interface

Description	This command enables RIB prioritization for the OSPF/OSPFv3 protocol. When enabled at the OSPF interface level, all routes learned through the associated OSPF interface will be processed through the OSPF route calculation process at a higher priority. The no form of rib-priority command disables RIB prioritization at the associated level.
Default	no rib-priority
Parameters	<i>prefix-list-name</i> — Specifies the prefix list which is used to select the routes that are processed at a higher priority through the route calculation process.

transit-delay

Syntax	transit-delay <i>seconds</i> no transit-delay
Context	config>service>vprn>ospf>area>if config>service>vprn>ospf3>area>if config>service>vprn>ospf>area>virtual-link config>service>vprn>ospf3>area>virtual-link config>service>vprn>ospf>area>sham-link
Description	This command configures the estimated time, in seconds, that it takes to transmit a link state advertisement (LSA) on the interface or virtual link or sham-link. The no form of this command reverts to the default delay time.
Default	transit-delay 1
Parameters	<i>seconds</i> — The transit delay in seconds expressed as a decimal integer. Values 0 to 3600

key-rollover-interval

Syntax	key-rollover-interval <i>key-rollover-interval</i>
Context	config>service>vprn>ospf3>area
Description	This command configures the key rollover interval. The no form of the command reverts to the default.
Default	key-rollover-interval 10
Parameters	<i>key-rollover-interval</i> — Specifies the time, in seconds, after which a key rollover will start. Values 10 to 300

loopfree-alternate-exclude

Syntax	[no] loopfree-alternate-exclude
Context	config>service>vprn>ospf>area config>service>vprn>ospf3>area
Description	This command specifies whether or not the OSPF/OSPF3 area should be excluded during LFA calculations. When enabled, the OSPF/OSPF3 area is excluded from LFA calculations. When disabled (the default), the OSPF/OSPF3 area is included in LFA calculations. The no form of the command includes the OSPF/OSPF3 area in LFA calculations.
Default	no loopfree-alternate-exclude

nssa

Syntax	[no] nssa
Context	config>service>vprn>ospf>area config>service>vprn>ospf3>area
Description	This command creates the context to configure an OSPF Not So Stubby Area (NSSA) and adds/removes the NSSA designation from the area. NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is an NSSA has the capability to flood external routes that it learns throughout its area and via an ABR to the entire OSPF domain. Existing virtual links of a non-stub or NSSA area will be removed when the designation is changed to NSSA or stub. An area can be designated as stub or NSSA but never both at the same time. By default, an area is not configured as an NSSA area. The no form of the command removes the NSSA designation and configuration context from the area.
Default	no nssa — The OSPF area is not an NSSA.

originate-default-route

Syntax	originate-default-route [type-7] [no-adjacency-check] no originate-default-route
Context	config>service>vprn>ospf>area>nssa

	config>service>vprn>ospf3>area>nssa
Description	This command enables the generation of a default route and its LSA type (3 or 7) into a Not So Stubby Area (NSSA) by an NSSA Area Border Router (ABR). When configuring an NSSA with no summaries, the ABR will inject a type 3 LSA default route into the NSSA area. Some older implementations expect a type 7 LSA default route. The no form of the command disables origination of a default route.
Default	no originate-default-route — A default route is not originated.
Parameters	type-7 — Specifies that a type 7 LSA should be used for the default route. Configure this parameter to inject a type 7 LSA default route into an NSSA configured with no summaries, instead of a type 3 LSA. To revert to a type 3 LSA, execute the originate-default-route command without the type-7 parameter. Default type 3 LSA default route no-adjacency-check — Specifies whether or not adjacency checks are performed before originating a default route. If this parameter is configured, then no area 0 adjacency is required for the ABR to advertise the default route. Default Adjacency checks are performed, and an area 0 adjacency is required for the ABR to advertise the default route

redistribute-external

Syntax	[no] redistribute-external
Context	config>service>vprn>ospf>area>nssa config>service>vprn>ospf3>area>nssa
Description	This command enables the redistribution of external routes into the Not So Stubby Area (NSSA) or an NSSA area border router (ABR) that is exporting the routes into non-NSSA areas. NSSA or Not So Stubby Areas are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is that the NSSA has the capability to flood external routes that it learns (providing it is an ASBR) throughout its area and via an Area Border Router to the entire OSPF domain. The no form of the command disables the default behavior to automatically redistribute external routes into the NSSA area from the NSSA ABR.
Default	redistribute-external — External routes are redistributed into the NSSA.

summaries

Syntax	[no] summaries
Context	config>service>vprn>ospf>area>nssa config>service>vprn>ospf>area>stub config>service>vprn>ospf3>area>nssa
Description	This command enables sending summary (type 3) advertisements into a stub area or Not So Stubby Area (NSSA) on an Area Border Router (ABR). This parameter is particularly useful to reduce the size of the routing and Link State Database (LSDB) tables within the stub or nssa area. By default, summary route advertisements are sent into the stub area or NSSA. The no form of the command disables sending summary route advertisements and, for stub areas, only the default route is advertised by the ABR.
Default	summaries — Summary routes are advertised by the ABR into the stub area or NSSA.

stub

Syntax	[no] stub
Context	config>service>vprn>ospf>area config>service>vprn>ospf3>area
Description	This command enables access to the context to configure an OSPF stub area and adds/removes the stub designation from the area. External routing information is not flooded into stub areas. All routers in the stub area must be configured with the stub command. An OSPF area cannot be both an NSSA and a stub area. Existing virtual links of a non STUB or NSSA area will be removed when its designation is changed to NSSA or STUB. By default, an area is not a stub area. The no form of the command removes the stub designation and configuration context from the area.
Default	no stub — The area is not configured as a stub area.

default-metric


Syntax	default-metric <i>metric</i> no default-metric
Context	config>service>vprn>ospf>area>stub config>service>vprn>ospf3>area>stub

Description	This command configures the metric used by the area border router (ABR) for the default route into a stub area. The default metric should only be configured on an ABR of a stub area. An ABR generates a default route if the area is a stub area. The no form of the command reverts to the default value.
Default	default-metric 1
Parameters	<i>metric</i> — The metric expressed as a decimal integer for the default route cost to be advertised into the stub area. Values 1 to 16777214

virtual-link

Syntax	[no] virtual-link <i>router-id</i> transit-area <i>area-id</i>
Context	config>service>vprn>ospf>area config>service>vprn>ospf3>area
Description	This command configures a virtual link to connect area border routers to the backbone via a virtual link. The backbone area (area 0.0.0.0) must be contiguous and all other areas must be connected to the backbone area. If it is not practical to connect an area to the backbone (see area 0.0.0.2 in Figure 53), then the area border routers (routers 1 and 2 in Figure 53) must be connected using a virtual link. The two area border routers will form a point-to-point like adjacency across the transit area (area 0.0.0.1 in Figure 53). A virtual link can only be configured while in the area 0.0.0.0 context. The <i>router-id</i> specified in this command must be associated with the virtual neighbor. The transit area cannot be a stub area or a Not So Stubby Area (NSSA). The no form of the command deletes the virtual link.
Default	No virtual link is defined.
Parameters	<i>router-id</i> — The router ID of the virtual neighbor in IP address dotted decimal notation. transit-area <i>area-id</i> — The area-id specified identifies the transit area that links the backbone area with the area that has no physical connection with the backbone. The OSPF backbone area, area 0.0.0.0, must be contiguous and all other areas must be connected to the backbone area. The backbone distributes routing information between areas. If it is not practical to connect an area to the backbone (see Area 0.0.0.5 in Figure 53) then the area border routers (such as routers Y and Z) must be connected via a virtual link. The two area border routers form a point-to-point-like adjacency across the transit area (see Area 0.0.0.4).

Figure 53 OSPF Areas



0991

compatible-rfc1583

Syntax [no] compatible-rfc1583

Context config>service>vprn>ospf

Description This command enables OSPF summary and external route calculations in compliance with RFC1583 and earlier RFCs.

RFC1583 and earlier RFCs use a different method to calculate summary and external route costs. To avoid routing loops, all routers in an OSPF domain should perform the same calculation method.

Although it would be favorable to require all routers to run a more current compliance level, this command allows the router to use obsolete methods of calculation.

This command is not supported in OSPF3.

The **no** form of the command enables the post-RFC1583 method of summary and external route calculation.

Default compatible-rfc1583 — RFC1583 compliance is enabled.

export

Syntax	export <i>policy-name</i> [<i>policy-name</i> ...] no export
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	This command associates export route policies to determine which routes are exported from the route table to OSPF. Export policies are only in effect if OSPF is configured as an ASBR. If no export policy is specified, non-OSPF routes are not exported from the routing table manager to OSPF. If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified. The no form of the command removes all policies from the configuration.
Default	no export — No export route policies specified.
Parameters	<i>policy-name</i> — The export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. The specified name(s) must already be defined.

external-db-overflow

Syntax	external-db-overflow <i>limit interval</i> no external-db-overflow
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	This command enables limits on the number of non-default AS-external-LSA entries that can be stored in the LSDB and specifies a wait timer before processing these after the limit is exceeded. The <i>limit</i> value specifies the maximum number of non-default AS-external-LSA entries that can be stored in the link-state database (LSDB). Placing a limit on the non-default AS-external-LSAs in the LSDB protects the router from receiving an excessive number of external routes that consume excessive memory or CPU resources. If the number of routes reach or exceed the <i>limit</i>, the table is in an overflow state. When in an overflow state, the router will not originate any new AS-external-LSAs. In fact, it withdraws all the self-originated non-default external LSAs.

The *interval* specifies the amount of time to wait after an overflow state before regenerating and processing non-default AS-external-LSAs. The waiting period acts like a dampening period preventing the router from continuously running Shortest Path First (SPF) calculations caused by the excessive number of non-default AS-external LSAs.

The **external-db-overflow** must be set identically on all routers attached to any regular OSPF area. OSPF stub areas and not-so-stubby areas (NSSAs) are excluded.

The **no** form of the command disables limiting the number of non-default AS-external-LSA entries.

Default	no external-db-overflow — No limit on non-default AS-external-LSA entries.
Parameters	<i>limit</i> — 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



Note: Setting a value of -1 is equivalent to **no external-db-overflow**.

interval — The number of seconds after entering an overflow state before attempting to process non-default AS-external-LSAs expressed as a decimal integer.

Values 0 to 2147483647

external-preference

Syntax	external-preference <i>preference</i> no external-preference
Context	config>service>vprn>ospf config>service>vprn>ospf3
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 per the default preference table as defined in the following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the ecmp in the config>router context.

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

Default external-preference 150 — OSPF external routes have a default preference of 150.

Parameters *preference* — The preference for external routes expressed as a decimal integer.

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes ¹
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

Note:

1. Preference for OSPF internal routes is configured with the **preference** command.

Values 1 to 255

ignore-dn-bit

Syntax **[no] ignore-dn-bit**

Context config>service>vprn>ospf

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.

loopfree-alternates

Syntax **[no] loopfree-alternates**

Context config>service>vprn>ospf
config>service>vprn>ospf3

Description This command enables Loop-Free Alternate (LFA) computation by SPF under the IS-IS routing protocol level, or under the OSPF routing protocol instance level.

When this command is enabled, it instructs the IGP SPF to attempt to pre-compute both a primary next-hop and an LFA next-hop for every learned prefix. IS-IS computes the primary SPF first and then computes the LFA SPF. The LFA backup next-hop is only available after the LFA SPF is completed. When found, the LFA next-hop is populated into the routing table along with the primary next-hop for the prefix.

The **no** form of this command disables the LFA computation by IGP SPF.

Default no loopfree-alternates

exclude

Syntax **exclude**

Context config>service>vprn>ospf>loopfree-alternates
config>service>vprn>ospf3>loopfree-alternates

Description This command excludes from LFA SPF calculation prefixes that match a prefix entry or a tag entry in a prefix policy.

The implementation already allows the user to exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.

If a prefix is excluded from LFA, then it will not be included in LFA calculation regardless of its priority. The prefix tag will, however, be used in the main SPF.



Note: Prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.

The default action of the **exclude** command, when not explicitly specified by the user in the prefix policy, is a “reject”. Thus, regardless if the user did or did not explicitly add the statement “default-action reject” to the prefix policy, a prefix that did not match any entry in the policy will be accepted into LFA SPF.

The **no** form of this command deletes the exclude prefix policy.

Default no exclude

prefix-policy

Syntax [**no**] **prefix-policy** *prefix-policy* [*prefix-policy*]

Context config>service>vprn>ospf>loopfree-alternates>exclude
config>service>vprn>ospf3>loopfree-alternates>exclude

Description	This command specifies the name of the policy for the prefixes to exclude from the LFA SPF calculation in this OSPF or OSPF3 instance. The no form of this command deletes the exclude prefix policy.
Default	no prefix-policy
Parameters	prefix-policy <i>prefix-policy</i> — Specifies the name of the prefix policy, up to 32 characters. Up to five prefix policies can be specified. The specified name must have been already defined.

overload

Syntax	overload [timeout seconds] no overload				
Context	config>service>vprn>ospf config>service>vprn>ospf3				
Description	This command changes the overload state of the local router so that it appears to be overloaded. When overload is enabled, the router can participate in OSPF routing, but is not used for transit traffic. Traffic destined to directly attached interfaces continue to reach the router. To put the IGP in an overload state enter a timeout value. The IGP will enter the overload state until the timeout timer expires or a no overload command is executed. If the overload command is encountered during the execution of an overload-on-boot command then this command takes precedence. This could occur as a result of a saved configuration file where both parameters are saved. When the file is saved by the system the overload-on-boot command is saved after the overload command. Use the no form of this command to return to the default. When the no overload command is executed, the overload state is terminated regardless the reason the protocol entered overload state.				
Default	no overload				
Parameters	timeout seconds — Specifies the number of seconds to reset overloading. <table> <tr> <td>Values</td><td>60 to 1800</td></tr> <tr> <td>Default</td><td>60</td></tr> </table>	Values	60 to 1800	Default	60
Values	60 to 1800				
Default	60				

if-attribute

Syntax	if-attribute
Context	config>router


```
config>router>interface
config>service>ies>interface
config>service>vprn>interface
```

Description This command creates the context to configure or apply IP interface attributes such as administrative group (admin-group) or Shared Risk Loss Group (SRLG).

admin-group

Syntax **admin-group** *group-name* [*group-name...*(up to 5 max)]
no admin-group *group-name* [*group-name...*(up to 5 max)]
no admin-group

Context config>router>if>if-attribute
config>service>ies>if>if-attribute
config>service>vprn>if>if-attribute
config>router>mpls>interface

Description This command configures the admin group membership of an interface. The user can apply admin groups to an IES, VPRN, network IP, or MPLS interface.

Each single operation of the **admin-group** command allows a maximum of five (5) groups to be specified at a time. However, a maximum of 32 groups can be added to a given interface through multiple operations. Once an admin group is bound to one or more interface, its value cannot be changed until all bindings are removed.

The configured admin-group membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.

Only the admin groups bound to an MPLS interface are advertised in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

The **no** form of this command deletes one or more of the admin-group memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.

Parameters *group-name* — Specifies the name of the group with up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain.

srlg-group

Syntax **srlg-group** *group-name* [*group-name...*(up to 5 max)]
no srlg-group *group-name* [*group-name...*(up to 5 max)]
no srlg-group

Context	<pre>config>router>if>if-attribute config>service>ies>if>if-attribute config>service>vprn>if>if-attribute config>router>mpls>interface</pre>
Description	This command configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface. An interface can belong to up to 64 SRLG groups. However, each single operation of the srlg-group command allows a maximum of five (5) groups to be specified at a time. Once an SRLG group is bound to one or more interface, its value cannot be changed until all bindings are removed. The configured SRLG membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas. Only the SRLGs bound to an MPLS interface are advertised in TE link TLVs and sub-TLVs when the traffic-engineering option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs. The no form of this command deletes one or more of the SRLG memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.
Parameters	<i>group-name</i> — Specifies the name of the group, up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain.

lfa-policy-map

Syntax	<pre>lfa-policy-map route-nh-template <i>template-name</i> no lfa-policy-map</pre>
Context	<pre>config>router>ospf>area>interface config>router>ospf3>area>interface config>router>isis>interface config>service>vprn>ospf>area>interface config>service>vprn>ospf3>area>interface</pre>
Description	This command applies a route next-hop policy template to an OSPF or IS-IS interface. When a route next-hop policy template is applied to an interface in IS-IS, it is applied in both level 1 and level 2. When a route next-hop policy template is applied to an interface in OSPF, it is applied in all areas. However, the command in an OSPF interface context can only be executed under the area in which the specified interface is primary and then applied in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command fails. If the user excluded the interface from LFA using the command loopfree-alternate-exclude, the LFA policy, if applied to the interface, has no effect.

Finally, if the user applied a route next-hop policy template to a loopback interface or to the system interface, the command will not be rejected, but it will result in no action being taken.

The **no** form deletes the mapping of a route next-hop policy template to an OSPF or IS-IS interface.

Parameters *template-name* — Specifies the name of the template, up to 32 characters.

overload-include-ext-2

Syntax	[no] overload-include-ext-2
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	This command is used to control if external type-2 routes should be re-advertised with a maximum metric value when the system goes into overload state for any reason. When this command is enabled and the router is in overload, all external type-2 routes will be advertised with the maximum metric.
Default	no overload-include-ext-2

overload-include-stub

Syntax	[no] overload-include-stub
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	This command is used to determine if the OSPF stub networks should be advertised with a maximum metric value when the system goes into overload state for any reason. When enabled, the system uses the maximum metric value. When this command is enabled and the router is in overload, all stub interfaces, including loopback and system interfaces, will be advertised at the maximum metric.
Default	no overload-include-stub

overload-on-boot

Syntax	overload-on-boot [timeout seconds] no overload
Context	config>service>vprn>ospf config>service>vprn>ospf3

Description	When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occur: • The timeout timer expires. • A manual override of the current overload state is entered with the no overload command. The no overload command does not affect the overload-on-boot function. The no form of the command removes the overload-on-boot functionality from the configuration.				
Default	no overload-on-boot				
Parameters	timeout <i>seconds</i> — Specifies the number of seconds to reset overloading. <table> <tr> <td>Values</td><td>60 to 1800</td></tr> <tr> <td>Default</td><td>60</td></tr> </table>	Values	60 to 1800	Default	60
Values	60 to 1800				
Default	60				

preference

Syntax	preference <i>preference</i> no preference
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	This command configures the preference for OSPF internal routes. A route can be learned by the router from different protocols in which case the costs are not comparable, when this occurs the preference is used to decide to which route will be used. Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the ecmp in the config>router context. The no form of the command reverts to the default value.
Default	preference 10 — OSPF internal routes have a preference of 10.
Parameters	<i>preference</i> — The preference for internal routes expressed as a decimal integer. Defaults for different route types are listed in the following table.

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes ¹
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

Note:

1. Preference for OSPF internal routes is configured with the **preference** command.

Values 1 to 255

reference-bandwidth

Syntax	reference-bandwidth <i>reference-bandwidth</i> reference-bandwidth [tbps <i>Tera-bps</i>] [gbps <i>Giga-bps</i>] [mbps <i>Mega-bps</i>] [kbps <i>Kilo-bps</i>] no reference-bandwidth
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	This command configures the reference bandwidth in kilobits per second (kb/s) that provides the reference for the default costing of interfaces based on their underlying link speed.

The default interface cost is calculated as follows:

$$\text{cost} = \text{reference-bandwidth} \div \text{bandwidth}$$

The default *reference-bandwidth* is 100,000,000 kb/s or 100 Gb/s, so the default auto-cost metrics for various link speeds are as follows:

- 10 Mb/s link default cost of 10000
- 100 Mb/s link default cost of 1000
- 1 Gb/s link default cost of 100
- 10 Gb/s link default cost of 10
- 40 Gb/s link default cost of 2
- 100 Gb/s link default cost of 1
- 400 Gb/s link default cost of 1



Warning: The default **reference-bandwidth** must be manually configured to a higher value, if interface speeds greater than 100 Gb/s and metrics based on link speed is used. When the default **reference-bandwidth** is used, a metric of 1 will be set on all interface speeds $\geq$ 100 Gb/s. For example, 100 GE, 100 GE LAG, 400 GE, and 400 GE LAG interfaces will all have a metric of 1.

The **reference-bandwidth** command assigns a default cost to the interface based on the interface speed. To override this default cost on a particular interface, use the **metric metric** command in the **config>router>ospf>area>interface ip-int-name** context.

The **no** form of the command reverts the reference-bandwidth to the default value.

Default	reference-bandwidth 100000000 — Reference bandwidth of 100 Gb/s.
Parameters	<i>reference-bandwidth</i> — The reference bandwidth in kilobits per second expressed as a decimal integer.
	Values 1 to 4000000000
	tbps <i>Tera-bps</i> — The reference bandwidth in terabits per second expressed as a decimal integer.
	Values 1 to 4
	gbps <i>Giga-bps</i> — The reference bandwidth in gigabits per second expressed as a decimal integer.
	Values 1 to 999
	mbps <i>Mega-bps</i> — The reference bandwidth in megabits per second expressed as a decimal integer.
	Values 1 to 999
	kbps <i>Kilo-bps</i> — reference bandwidth in kilobits per second expressed as a decimal integer.
	Values 1 to 999

rib-priority

Syntax	rib-priority {high} prefix-list-name no rib-priority
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	This command enabled RIB prioritization for the OSPF protocol and specifies the prefix list that will be used to select the specific routes that should be processed through the OSPF route calculation process at a higher priority.

The **no** form of **rib-priority** command disables RIB prioritization at the associated level.

- Default** no rib-priority
- Parameters** *prefix-list-name* — Specifies the prefix list which is used to select the routes that are processed at a higher priority through the route calculation process.

rtr-adv-lsa-limit

- Syntax** **rtr-adv-lsa-limit** [1..4294967295] [**log-only**] [**threshold** *percent*]
rtr-adv-lsa-limit [1..4294967295] [**log-only**] [**threshold** *percent*] **overload-timeout forever**
rtr-adv-lsa-limit [1..4294967295] [**log-only**] [**threshold** *percent*] **overload-timeout** *seconds*
no rtr-adv-lsa-limit
- Context** config>service>vprn>ospf
- Description** This command configures the maximum number of LSAs OSPF can learn from another router, in order to protect the system from a router that accidentally advertises a large number of LSAs. When the number of advertised LSAs reaches the configured percentage of this limit, an SNMP trap is sent. If the limit is exceeded, OSPF goes into overload.
- The **overload-timeout** option allows the administrator to control how long OSPF is in overload as a result of the advertised LSA limit being reached. At the end of this duration of time the system automatically attempts to restart OSPF. One possible value for the **overload-timeout** is **forever**, which means OSPF is never restarted automatically and this corresponds to the default behavior when the **overload-timeout** option is not configured.
- The **no** form of the command removes the **rtr-adv-lsa-limit**.
- Default** forever
- Parameters** **log-only** — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, overload is not set.
- percent* — The threshold value (as a percentage) that triggers a warning message to be sent.
- Values** 0 to 100
- seconds* — Specifies duration in seconds before restarting OSPF.
- Values** 1 to 1800

super-backbone

- Syntax** [**no**] **super-backbone**
- Context** config>service>vprn>ospf

Description	This command specifies whether CE-PE functionality is required or not. The OSPF super backbone indicates the type of the LSA generated as a result of routes redistributed into OSPF. When enabled, the redistributed routes are injected as summary, external or NSSA LSAs. When disabled, the redistributed routes are injected as either external or NSSA LSAs only.
Default	no super-backbone

suppress-dn-bit

Syntax	[no] suppress-dn-bit
Context	config>service>vprn>ospf config>service>vprn>ospf3
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 config>service>vprn>ospf3
Description	This command enters the context that allows for the configuration of OSPF timers. Timers control the delay between receipt of a link state advertisement (LSA) requiring a Dijkstra (Shortest Path First (SPF)) calculation and the minimum time between successive SPF calculations. Changing the timers affect CPU utilization and network reconvergence times. Lower values reduce convergence time but increase CPU utilization. Higher values reduce CPU utilization but increase reconvergence time.
Default	n/a

incremental-spf-wait

Syntax	incremental-spf-wait <i>inc-spf-wait</i> no incremental-spf-wait

Context	config>router>ospf>timers config>router>ospf3>timers
Description	This command sets the delay before an incremental SPF calculation is performed when LSA types 3, 4, 5, or 7 are received. This allows multiple updates to be processed in the same SPF calculation. Type 1 or type 2 LSAs are considered a topology change and will always trigger a full SPF calculation. The no incremental-spf-wait form of the command resets the timer value back to the default value.
	Note: The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is ≥ 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.
Default	incremental-spf-wait 1000
Parameters	<i>inc-spf-wait</i> — Specifies the OSPF incremental SPF calculation delay, in milliseconds.
	Values 0 to 1000

lsa-accumulate

Syntax	lsa-accumulate <i>lsa-accum-time</i> no lsa-accumulate
Context	config>router>ospf>timers config>router>ospf3>timers
Description	This commands sets the internal OSPF delay to allow for the accumulation of multiple LSA so OSPF messages can be sent as efficiently as possible. The lsa-accumulate timer applies to all LSAs, except Type 1 and Type 2 LSAs which are sent immediately. LSAs are accumulated and then sent when: • Its size reaches the MTU size of the interface. • A new LSA is received on the interface. • The lsa-accumulate timer expires. Shorting this delay can speed up the advertisement of LSAs to OSPF neighbors but may increase the number of OSPF messages sent.
	Note: The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is ≥ 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Default	<code>lsa-accumulate 1000</code>
Parameters	<i>lsa-accum-time</i> — Specifies the LSA accumulation delay in milliseconds.
Values	0 to 1000

lsa-arrival

Syntax	lsa-arrival <i>lsa-arrival-time</i> no lsa-arrival
Context	config>service>vprn>ospf>timers config>service>vprn>ospf3>timers
Description	This parameter defines the minimum delay that must pass between receipt of the same Link State Advertisements (LSAs) arriving from neighbors. It is recommended that the neighbor's configured lsa-generate <i>lsa-second-wait</i> interval is equal to or greater than the lsa-arrival timer configured here. Use the no form of this command to return to the default.



Note: The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is ≥ 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Default	<code>lsa-arrival 1000</code>
Parameters	<i>lsa-arrival-time</i> — Specifies the timer in milliseconds.
Values	0 to 600000

lsa-generate

Syntax	lsa-generate <i>max-lsa-wait</i> [<i>lsa-initial-wait lsa-initial-wait [lsa-second-wait lsa-second-wait]</i>] no lsa-generate-interval
Context	config>service>vprn>ospf>timers config>service>vprn>ospf3>timers
Description	This parameter customizes the throttling of OSPF LSA-generation. Timers that determine when to generate the first, second, and subsequent LSAs can be controlled with this command. Subsequent LSAs are generated at increasing intervals of the <i>lsa-second-wait</i> timer until a maximum value is reached. Configuring the lsa-arrival interval to equal or less than the <i>lsa-second-wait</i> interval configured in the lsa-generate command is recommended.

Use the **no** form of this command to return to the default.



Note: The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is ≥ 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Parameters	<i>max-lsa-wait</i> — Specifies the maximum interval, in milliseconds, between two consecutive occurrences of an LSA being generated.
	Values 10 to 600000
	Default 5000
	<i>lsa-initial-wait</i> — Specifies the first waiting period between link-state advertisements (LSA) originate(s), in milliseconds. When the LSA exceeds the lsa-initial-wait timer value and the topology changes, there is no wait period and the LSA is immediately generated.
	When an LSA is generated, the initial wait period commences. If, within the specified lsa-initial-wait period and another topology change occurs, then the lsa-initial-wait timer applies.
	Values 10 to 600000
	Default 5000
	<i>lsa-second-wait</i> — Specifies the hold time in milliseconds between the first and second LSA generation. The next topology change is subject to this second wait period. With each subsequent topology change, the wait time doubles (this is 2x the previous wait time). This assumes that each failure occurs within the relevant wait period.
	Values 10 to 600000
	Default 5000

redistribute-delay

Syntax	redistribute-delay <i>redist-wait</i> no redistribute-delay
Context	config>router>ospf>timers config>router>ospf3>timers
Description	This command sets the internal OSPF hold down timer for external routes being redistributed into OSPF. Shorting this delay can speed up the advertisement of external routes into OSPF but can result in additional OSPF messages if that source route is not yet stable. The no redistribute-delay form of the command resets the timer value back to the default value.



Note: The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is ≥ 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Default	redistribute-delay 1000
Parameters	<i>redist-wait</i> — Specifies the OSPF redistribution hold down timer, in milliseconds, for external routes being advertised into OSPF.
Values	0 to 1000

spf-wait

Syntax	spf-wait <i>max-spf-wait</i> [<i>spf-initial-wait</i> <i>spf-initial-wait</i>] [<i>spf-second-wait</i> <i>spf-second-wait</i>] no spf-wait
Context	config>service>vprn>ospf>timers config>service>vprn>ospf3>timers
Description	This command defines the maximum interval between two consecutive SPF calculations in milliseconds. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the <i>spf-second-wait</i> interval. For example, if the <i>spf-second-wait</i> interval is 1000, then the next SPF will run after 2000 milliseconds, and then 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 <i>spf-initial-wait</i> .

Use the **no** form of this command to return to the default.



Note: The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is ≥ 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

Parameters	<i>max-spf-wait</i> — Specifies the maximum interval in milliseconds between two consecutive SPF calculations.
Values	10 to 120000
Default	10000
	<i>spf-initial-wait</i> — 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

unicast-import-disable

Syntax **[no] unicast-import-disable**

Context config>service>vprn>ospf

Description This command allows one IGP to import its routes into RPF RTM while another IGP imports routes only into the unicast RTM.

Import policies can redistribute routes from an IGP protocol into the RPF RTM (the multicast routing table). By default, the IGP routes will not be imported into RPF RTM as such an import policy must be explicitly configured

Default no unicast-import-disable

vpn-domain

Syntax **vpn-domain** [*type* {0005 | 0105 | 0205 | 8005}] **id** *id*
no vpn-domain

Context config>service>vprn>ospf

Description This command specifies type of the extended community attribute exchanged using BGP to carry the OSPF VPN domain ID. This applies to VPRN instances of OSPF only. An attempt to modify the value of this object will result in an inconsistent value error when is not a VPRN instance. The parameters are mandatory and can be entered in either order. This command is not applicable in the **config>service>vprn>ospf3** context.

This command is not supported in OSPF3.

Default no vpn-domain

Parameters *id* — Specifies the OSPF VPN domain in the “xxxx.xxxx.xxxx” format. This is exchanged using BGP in the extended community attribute associated with a prefix. This object applies to VPRN instances of OSPF only.

type — Specifies the type of the extended community attribute exchanged using BGP to carry the OSPF VPN domain ID.

Values 0005, 0105, 0205, 8005

vpn-tag

Syntax	vpn-tag <i>vpn-tag</i> no vpn-tag
Context	config>service>vprn>ospf
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 will result in an inconsistent value error when is not a VPRN instance. This command is not supported in OSPF3.
Default	vpn-tag 0

3.8.2.32 PIM Commands

pim

Syntax	[no] pim
Context	config>service>vprn
Description	This command configures a Protocol Independent Multicast (PIM) instance in the VPRN service. When an PIM instance is created, the protocol is enabled. PIM is used for multicast routing within the network. Devices in the network can receive the multicast feed requested and non-participating routers can be pruned. The router supports PIM sparse mode (PIM-SM). The no form of the command deletes the PIM protocol instance removing all associated configuration parameters.

apply-to

Syntax	apply-to {all none}
Context	config>service>vprn>pim
Description	This command creates a PIM interface with default parameters. If a manually created interface or modified interface is deleted, the interface will be recreated when the apply-to command is executed. If PIM is not required on a specific interface, then execute a shutdown command.

The **apply-to** command is saved first in the PIM configuration structure, all subsequent commands either create new structures or modify the defaults as created by the **apply-to** command.

Default	apply-to none
Parameters	all — Specifies that all VPRN and non-VPRN interfaces are automatically applied in PIM. none — No interfaces are automatically applied in PIM. PIM interfaces must be manually configured.

grt-extranet

Syntax	[no] grt-extranet
Context	config>service>vprn>pim
Description	This command enters the context to configure GRT/VRF extranet for this MVPN instance.

group-prefix

Syntax	group-prefix <i>ip-address/mask</i> [<i>ip-address/mask...(up to 8 max)</i>] [starg] group-prefix any no group-prefix <i>ip-address/mask</i> no group-prefix any
Context	config>service>vprn>pim>rpf-select>grt-extranet
Description	This command configures multicast group IPv4 prefixes for the multicast GRT/VRF with per group mapping extranet functionality. Multiple lines are allowed. Duplicate prefixes are ignored. Operator can either configure specific groups for extranet or specify all groups by using key-word any. The two options are mutually exclusive in configuration. When the starg option is specified, extranet functionality is enabled for PIM ASM as for the specified group. When the option is not specified (not recommended with PIM ASM), the PIM ASM join will be mapped and data plane will be established, but the control plane will not be updated on SPT switchover, unless the switchover is driven by a CPE router on a receiver side. The no form of the command deletes specified prefix from the list, or removes mapping of all prefixes if group-prefix any was specified.
Parameters	<i>ip-address/mask</i> — Specifies the IPv4 multicast address prefix with mask.

import

Syntax	import { join-policy register-policy } <i>policy-name</i> [<i>policy-name</i> ...(up to 5 max)] no import { join-policy register-policy }
Context	config>service>vprn>pim
Description	This command specifies the import route policy to be used for determining which routes are accepted from peers. Route policies are configured in the config>router>policy-options context. When an import policy is not specified, BGP routes are accepted by default. The no form of the command removes the policy association from the IGMP instance.
Default	no import join-policy no import register-policy
Parameters	join-policy — Use this command to filter PIM join messages which prevents unwanted multicast streams from traversing the network. register-policy — This keyword filters register messages. PIM register filters prevent register messages from being processed by the RP. This filter can only be defined on an RP. When a match is found, the RP immediately sends back a register-stop message. <i>policy-name</i> — The route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. Route policies are configured in the config>router>policy-options context.

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>service>vprn>pim
Description	This command enables PIM on an interface and enables the context to configure interface-specific parameters. By default interfaces are activated in PIM based on the apply-to command, and do not have to be configured on an individual basis unless the default values must be changed. The no form of the command deletes the PIM interface configuration for this interface. If the apply-to command parameter is configured, then the no interface form must be saved in the configuration to avoid automatic (re)creation after the next apply-to is executed as part of a reboot. The shutdown command can be used to disable an interface without removing the configuration for the interface.

Default	Interfaces are activated in PIM based on the apply-to command.
Parameters	<i>ip-int-name</i> — Specifies the interface name. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

assert-period

Syntax	assert-period <i>assert-period</i> no assert-period
Context	config>service>vprn>pim>if
Description	This command configures the period in seconds for periodic refreshes of PIM Assert messages on an interface. The no form of the command reverts to the default.
Default	60
Parameters	<i>assert-period</i> — Specifies the period, in seconds, for periodic refreshes of PIM Assert messages on an interface. Values 1 to 300

bfd-enable

Syntax	[no] bfd-enable [ipv4 ipv6]
Context	config>service>vprn>pim>if
Description	This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface. The no form of this command removes BFD from the associated IGP protocol adjacency.
Default	no bfd-enable

bsm-check-rtr-alert

Syntax	[no] bsm-check-rtr-alert
Context	config>service>vprn>pim>if
Description	This command enables the checking of router alert option in the bootstrap messages received on this interface.

Default no bsm-check-rtr-alert

hello-interval

Syntax **hello-interval** *hello-interval*
 no hello-interval

Context config>service>vprn>pim>if

Description This command configures the frequency at which PIM Hello messages are transmitted on this interface.

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

Default hello-interval 30

Parameters *hello-interval* — Specifies the hello interval in seconds. A 0 (zero) value disables the sending of Hello messages (the PIM neighbor will never timeout the adjacency).

Values 0 to 255 seconds

hello-multiplier

Syntax **hello-multiplier** *deci-units*
 no hello-multiplier

Context config>service>vprn>pim>if

Description This command configures the multiplier to determine the holdtime for a PIM neighbor on this interface.

 The **hello-multiplier** in conjunction with the **hello-interval** determines the holdtime for a PIM neighbor.

Default hello-multiplier 35

Parameters *deci-units* — Specify the value, specified in multiples of 0.1, for the formula used to calculate the holdtime based on the **hello-multiplier**:

 (**hello-interval** * **hello-multiplier**) / 10

 This allows the PIMv2 default **hello-multiplier** of 3.5 and the default timeout of 105 seconds to be supported.

Values 20 to 100

improved-assert

Syntax [no] improved-assert

Context	config>service>vprn>pim>if
Description	This command enables improved assert processing on this interface. The PIM assert process establishes a forwarder for a LAN and requires interaction between the control and forwarding planes. The assert process is started when data is received on an outgoing interface. This could impact performance if data is continuously received on an outgoing interface. When enabled, the PIM assert process is done entirely on the control-plane with no interaction between the control and forwarding plane.
Default	improved-assert

instant-prune-echo

Syntax	[no] instant-prune-echo
Context	config>service>vprn>pim>interface
Description	This command enables PIM to send an instant prune echo when the router starts the prune pending timer for a group on the interface. All downstream routers will see the prune message immediately, and can send a join override if they are interested in receiving the group. Configuring instant-prune-echo is recommended on broadcast interfaces with more than one PIM neighbor to optimize multicast convergence. The no form of the command disables instant Prune Echo on the PIM interface.
Default	no instant-prune-echo

max-groups

Syntax	max-groups <i>value</i> no max-groups
Context	config>service>vprn>pim>if
Description	This command configures the maximum number of groups for which PIM can have downstream state based on received PIM Joins on this interface. This does not include IGMP local receivers on the interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed. When this object has a value of 0, there is no limit to the number of groups.
Parameters	<i>value</i> — Specifies the maximum number of groups for this interface.
Values	1 to 16000

monitor-oper-group

Syntax	monitor-oper-group <i>group-name</i> family { ipv4 ipv6 } { add set subtract } <i>value</i> no monitor-oper-group [family { ipv4 ipv6 }]
Context	config>service>vprn>pim>if
Description	This command configures PIM to monitor the state of an operational group to provide a redundancy mechanism. PIM monitors the operational group and changes its DR priority to the specified value when the status of the operational group is up. This enables the router to become the PIM DR only when the operational group is up. If the operational group status changes to down, PIM changes its DR priority to the default or the value configured with priority under config>service>vprn>pim>if. The oper-group <i>group-name</i> must already be configured under the config>service context before its name is referenced in this command. Two operational groups are supported per PIM interface. The no form of the command removes the operational group from the configuration.
Parameters	<i>group-name</i> — Specifies the operational group identifier up to 32 characters in length. family — Specifies the address family. ipv4 — Specifies the IPv4 designated router priority. ipv6 — Specifies the IPv6 designated router priority. add — Specifies that the value is to be added to the existing priority to become the designated router. subtract — Specifies that the value is to be subtracted from the existing priority to become the designated router. set — Specifies the priority to become the designated router. <i>value</i> — Specifies the priority modifier expressed as a decimal integer. Values 1 to 4294967295

multicast-senders

Syntax	multicast-senders { auto always never } no multicast-senders
Context	config>service>vprn>pim>if
Description	This command configures the way subnet matching is done for incoming data packets on this interface. An IP multicast sender is an user entity to be authenticated in a receiving host.
Parameters	auto — Subnet matching is automatically performed for incoming data packets on this interface. always — Subnet matching is always performed for incoming data packets on this interface.

never — Subnet matching is never performed for incoming data packets on this interface.

p2mp-ldp-tree-join

Syntax	p2mp-ldp-tree-join p2mp-ldp-tree-join ipv4 p2mp-ldp-tree-join ipv6 p2mp-ldp-tree-join ipv4 ipv6 no p2mp-ldp-tree-join [ipv4] [ipv6]
Context	config>service>vprn>pim>if
Description	This command configures the option to join P2MP LDP tree towards the multicast source for the VPRN service. If p2mp-ldp-tree-join is enabled, a PIM multicast join received on an interface is processed to join P2MP LDP LSP using the in-band signaled P2MP tree for the same multicast flow. LDP P2MP tree is setup towards the multicast source. Route to source of the multicast node is looked up from the RTM. The next-hop address for the route to source is set as the root of LDP P2MP tree. The no form of command disables joining P2MP LDP tree for IPv4 or IPv6 or both (if both or none is specified).
Default	no p2mp-ldp-tree-join
Parameters	<i>ipv4</i> — Enables dynamic mLDLP in-band signaling for IPv4 PIM joins. IPv4 multicast must be enabled; see ipv4-multicast-disable. For backward compatibility p2mp-ldp-tree-join is equivalent to p2mp-ldp-tree-join ipv4. <i>ipv6</i> — Enables dynamic mLDLP in-band signaling for IPv6 PIM joins. IPv6 multicast must be enabled; see ipv6-multicast-disable.

priority

Syntax	priority dr-priority no priority
Context	config>service>vprn>pim>if
Description	This command sets the priority value to become the rendezvous point (RP) that is included in bootstrap messages sent by the router. The RP is sometimes called the bootstrap router. The priority command indicates whether the router is eligible to be a bootstrap router. The no form of the command disqualifies the router to participate in the bootstrap election.
Default	priority 1 (The router is the least likely to become the designated router.)

Parameters	<i>dr-priority</i> — Specifies the priority to become the designated router. The higher the value, the higher the priority.
Values	1 to 4294967295

sticky-dr

Syntax	sticky-dr [<i>priority dr-priority</i>] no sticky-dr
Context	config>service>vprn>pim>if
Description	This command enables sticky-dr operation on this interface. When enabled, the priority in PIM hellos sent on this interface when elected as the designated router (DR) will be modified to the value configured in <i>dr-priority</i>. This is done to avoid the delays in forwarding caused by DR recovery, when switching back to the old DR on a LAN when it comes back up. By enabling sticky-dr on this interface, it will continue to act as the DR for the LAN even after the old DR comes back up. The no form of the command disables sticky-dr operation on this interface.
Default	no sticky-dr
Parameters	priority <i>dr-priority</i> — Sets the DR priority to be sent in PIM Hello messages following the election of that interface as the DR, when sticky-dr operation is enabled.
Values	1 to 4294967295

three-way-hello

Syntax	three-way-hello [<i>compatibility-mode</i>] no three-way-hello
Context	config>service>vprn>pim>if
Description	This command configures the compatibility mode for enabling the three way hello.
Parameters	compatibility-mode — Specifies to enable the three way hello.

tracking-support

Syntax	[no] tracking-support
Context	config>service>vprn>pim>if
Description	This command sets the T bit in the LAN Prune Delay option of the Hello Message. This indicates the router's capability to disable Join message suppression.

Default no tracking-support

ipv4-multicast-disable

Syntax [no] **ipv4-multicast-disable**

Context config>service>vprn>pim
config>service>vprn>pim>interface

Description This command administratively disables/enables PIM operation for IPv4.

Default no ipv4-multicast-disable

ipv6-multicast-disable

Syntax **ipv6-multicast-disable**

Context config>service>vprn>pim
config>service>vprn>pim>interface

Description This command administratively disables/enables PIM operation for IPv6.

Default ipv6-multicast-disable (config>service>vprn>pim)
no ipv6-multicast-disable (config>service>vprn>pim>interface)

mc-ecmp-balance

Syntax [no] **mc-ecmp-balance**

Context config>service>vprn>pim

Description This command enables multicast balancing of traffic over ECMP links based on the number of (S, G) distributed over each link. When enabled, each new multicast stream that needs to be forwarded over an ECMP link is compared to the count of (S, G) already using each link, so that the link with the fewest (S, G) is chosen.

This command cannot be used together with the **mc-ecmp-hashing-enabled** command.

The **no** form of the command disables multicast ECMP balancing.

mc-ecmp-balance-hold

Syntax **mc-ecmp-balance-hold** *minutes*
no mc-ecmp-balance-hold

Context	config>service>vprn>pim
Description	This command configures the hold time for multicast balancing over ECMP links.
Parameters	<i>minutes</i> — Specifies the hold time, in minutes, that applies after an interface has been added to the ECMP link.

mc-ecmp-hashing-enabled

Syntax	mc-ecmp-hashing-enabled [rebalance] no mc-ecmp-hashing-enabled
Context	config>service>vprn>pim
Description	This command enables hash-based multicast balancing of traffic over ECMP links and causes PIM joins to be distributed over the multiple ECMP paths based on a hash of S and G (and possibly next-hop IP address). When a link in the ECMP set is removed, the multicast flows that were using that link are redistributed over the remaining ECMP links using the same hash algorithm. When a link is added to the ECMP set new joins may be allocated to the new link based on the hash algorithm, but existing multicast flows using the other ECMP links stay on those links until they are pruned. Hash-based multicast balancing is supported for both IPv4 and IPv6. This command cannot be used together with the mc-ecmp-balance command. Using this command and the lag-usage-optimization command on mixed port speed LAGs is not recommended, because some groups may be forwarded incorrectly. The no form of the command disables the hash-based multicast balancing of traffic over ECMP links. The no mc-ecmp-hashing-enabled form of the command means that the use of multiple ECMP paths if enabled at the config>router or config>service>vprn context is controlled by the existing implementation and CLI commands mc-ecmp-balance.
Default	no mc-ecmp-hashing-enabled
Parameters	rebalance — Specifies to rebalance flows to newly added links immediately, instead of waiting until they are pruned.

non-dr-attract-traffic

Syntax	[no] non-dr-attract-traffic
Context	config>service>vprn>pim
Description	This command specifies whether the router should ignore the designated router state and attract traffic even when it is not the designated router.

An operator can configure an interface (router or IES or VPRN interfaces) to IGMP and PIM. The interface IGMP state will be synchronized to the backup node if it is associated with the redundant peer port. The interface can be configured to use PIM which will cause multicast streams to be sent to the elected DR only. The DR will also be the router sending traffic to the DSLAM. Since it may be required to attract traffic to both routers a flag non-dr-attract-traffic can be used in the PIM context to have the router ignore the DR state and attract traffic when not DR. While using this flag, the router may not send the stream down to the DSLAM while not DR.

When enabled, the designated router state is ignored. When disabled, **no non-dr-attract-traffic**, the designated router value is honored.

Default no non-dr-attract-traffic

rpf6-table

Syntax **rpf6-table {rtable6-m | rtable6-u | both}**
no rpf6-table

Context config>service>vprn>pim

Description This command configures the sequence of route tables used to find a Reverse Path Forwarding (RPF) interface for a specific multicast route.

By default, only the unicast route table is looked up to calculate the RPF interface toward the source/rendezvous point. However, the operator can specify to use the following:

- unicast route table only
- multicast route table only
- both route tables

Default rpf6-table rtable6-u

Parameters **rtable6-m** — Specifies that only the multicast route table will be used by the multicast protocol (PIM) for IPv6 RPF checks. This route table will contain routes submitted by static routes, ISIS and OSPF.

rtable6-u — Specifies that only the unicast route table will be used by the multicast protocol (PIM) for IPv6 RPF checks. This route table will contain routes submitted by all unicast routing protocols.

both — Specifies that the multicast route table will be used first by the multicast protocol (PIM) for IPv6 RPF checks, then the unicast route table will be used if the multicast route table lookup fails.

rp

Syntax **rp**

Context	config>service>vprn>pim
Description	This command enables access to the context to configure the rendezvous point (RP) of a PIM protocol instance. A Nokia PIM router acting as an RP must respond to a PIM register message specifying an SSM multicast group address by sending stop register message(s) to the first hop router. It does not build an (S, G) shortest path tree toward the first hop router. An SSM multicast group address can be either from the SSM default range of 232/8 or from a multicast group address range that was explicitly configured for SSM.
Default	rp enabled when PIM is enabled.

anycast

Syntax	[no] anycast <i>rp-ip-address</i>
Context	config>service>vprn>pim>rp
Description	This command configures a PIM anycast protocol instance for the RP being configured. Anycast enables fast convergence when a PIM RP router fails by allowing receivers and sources to rendezvous at the closest RP. The no form of the command removes the anycast instance from the configuration.
Parameters	<i>rp-ip-address</i> — Configure the loopback IP address shared by all routes that form the RP set for this anycast instance. Only a single address can be configured. If another anycast command is entered with an address then the old address will be replaced with the new address. If no ip-address is entered then the command is simply used to enter the anycast CLI level. Values Any valid loopback address configured on the node.

rp-set-peer

Syntax	[no] rp-set-peer <i>ip-address</i>
Context	config>service>vprn>pim>rp>anycast
Description	This command configures a peer in the anycast RP-set. The address identifies the address used by the other node as the RP candidate address for the same multicast group address range as configured on this node. This is a manual procedure. Caution should be taken to produce a consistent configuration of an RP-set for a given multicast group address range. The priority should be identical on each node and be a higher value than any other configured RP candidate that is not a member of this RP-set.

Although there is no set maximum of addresses that can be configured in an RP-set, up to 15 multicast addresses is recommended.

The **no** form of the command removes an entry from the list.

Parameters *ip-address* — Specifies the address used by the other node as the RP candidate address for the same multicast group address range as configured on this node.

auto-rp-discovery

Syntax **[no] auto-rp-discovery**

Context config>service>vprn>pim>rp

Description This command enables auto-RP protocol in discovery mode. In discovery mode, RP-mapping and RP-candidate messages are received and forwarded to downstream nodes. RP-mapping messages are received locally to learn about availability of RP nodes present in the network. In a VPRN configuration, it is recommended that a local loopback interface be created with the same IP address as the system IP address.

Either **bsr-candidate** for IPv4 or **auto-rp-discovery** can be configured; the two mechanisms cannot be enabled together. **bsr-candidate** for IPv6 and **auto-rp-discovery** for IPv4 can be enabled together. **auto-rp-discovery** cannot be enabled together with **mdt-type sender-only** or **mdt-type receiver-only**, or **wildcard-spmsi** configurations.

The **no** form of the command disables auto-RP.

Default no auto-rp-discovery

bootstrap-export

Syntax **bootstrap-export** *policy-name* [*policy-name... up to five*]
no bootstrap-export

Context config>service>vprn>pim>rp

Description This command exports policies to control the flow of bootstrap messages from the RP. Up to five policies can be defined.

The **no** form of this command removes the specified policy names from the configuration.

Parameters *policy-name* — Specifies the policy name. The policy statement must already be configured in the config>router>policy-options context.

bootstrap-import

Syntax **bootstrap-import** *policy-name* [*policy-name... up to five*]

no bootstrap-import *policy-name* [*policy-name*... up to five]

Context	config>service>vprn>pim>rp
Description	This command imports policies to control the flow of bootstrap messages into the RP. Up to five policies can be defined. The no form of this command removes the specified policy names from the configuration.
Parameters	<i>policy-name</i> — Specifies the policy name. The policy statement must already be configured in the config>router>policy-options context.

bsr-candidate

Syntax	bsr-candidate
Context	config>service>vprn>pim>rp config>service>vprn>pim>rp>ipv6
Description	This command enters the context to configure Candidate Bootstrap (BSR) parameters. Either bsr-candidate for IPv4 or auto-rp-discovery can be configured; the two mechanisms cannot be enabled together. bsr-candidate for IPv6 and auto-rp-discovery for IPv4 can be enabled together. The no form of the command disables BSR.
Default	no bsr-candidate

address

Syntax	[no] address <i>ip-address</i>
Context	config>service>vprn>pim>rp>bsr-candidate config>service>vprn>pim>rp>rp-candidate
Description	This command configures a static bootstrap or rendezvous point (RP) as long as the source is not directly attached to this router. Use the no form of this command to remove the static RP from the configuration.
Default	No IP address is specified.
Parameters	<i>ip-address</i> — The static IP address of the RP. The <i>ip-address</i> portion of the address command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Values 1.0.0.0 to 223.255.255.255

address

Syntax	[no] address <i>ipv6-address</i>
Context	config>service>vprn>pim>rp>ipv6>bsr-candidate config>service>vprn>pim>rp>ipv6>rp-candidate
Description	This command configures a static bootstrap or rendezvous point (RP) as long as the source is not directly attached to this router. Use the no form of this command to remove the static RP from the configuration.
Default	No IP address is specified.
Parameters	<i>ipv6-address</i> — The static IP address of the RP. The <i>ip-address</i> 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 <i>ipv6-address</i> : x: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

hash-mask-len

Syntax	hash-mask-len <i>hash-mask-length</i> no hash-mask-len
Context	config>service>vprn>pim>rp>bsr-candidate
Description	This command is used to configure the length of a mask that is to be combined with the group address before the hash function is called. All groups with the same hash map to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This mechanism is used to map one group or multiple groups to an RP.
Default	hash-mask-len 30
Parameters	<i>hash-mask-length</i> — The hash mask length. Values 0 to 32

hash-mask-len

Syntax	hash-mask-len <i>hash-mask-length</i> no hash-mask-len

Context	config>service>vprn>pim>rp>ipv6>bsr-candidate
Description	This command is used to configure the length of a mask that is to be combined with the group address before the hash function is called. All groups with the same hash map to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This mechanism is used to map one group or multiple groups to an RP.
Default	hash-mask-len 126
Parameters	<i>hash-mask-length</i> — The hash mask length.
Values	0 to 128

priority

Syntax	priority <i>bootstrap-priority</i>
Context	config>service>vprn>pim>rp>bsr-candidate config>service>vprn>pim>rp>ipv6>bsr-candidate
Description	This command defines the priority used to become the rendezvous point (RP). The higher the priority value the more likely that this router becomes the RP. If there is a tie, the router with the highest IP address is elected.
Parameters	<i>bootstrap-priority</i> — The priority to become the bootstrap router.
Values	0 to 255
Default	0 (the router is not eligible to be the bootstrap router)

ipv6

Syntax	ipv6
Context	config>service>vprn>pim>rp
Description	This command enables access to the context to configure the rendezvous point (RP) of a PIM IPv6 protocol instance. A Nokia IPv6 PIM router acting as an RP must respond to an IPv6 PIM register message specifying an SSM multicast group address by sending to the first hop router stop register message(s). It does not build an (S, G) shortest path tree toward the first hop router. An SSM multicast group address can be either from the SSM default range or from a multicast group address range that was explicitly configured for SSM.
Default	ipv6 RP enabled when IPv6 PIM is enabled.

anycast

Syntax	anycast <i>ipv6-address</i> no anycast <i>ipv6-address</i>
Context	config>service>vprn>pim>rp>ipv6
Description	This command configures an IPv6 PIM anycast protocol instance for the RP being configured. Anycast enables fast convergence when a PIM RP router fails by allowing receivers and sources to rendezvous at the closest RP. The no form of the command removes the anycast instance from the configuration.
Parameters	<i>ipv6-address</i> — Configures the loopback IP address shared by all routes that form the RP set for this anycast instance. Only a single address can be configured. If another anycast command is entered with an address then the old address will be replaced with the new address. If no address is entered then the command is simply used to enter the anycast CLI context. Values <i>ipv6-address</i> : 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

rp-set-peer

Syntax	[no] rp-set-peer <i>ipv6-address</i>
Context	config>service>vprn>pim>rp>ipv6>anycast
Description	This command configures an IPv6 peer in the anycast rp-set. The address identifies the address used by the other node as the RP candidacy address for the same multicast group address range as configured on this node. This is a manual procedure. Caution should be taken to produce a consistent configuration of an RP- set for a given multicast group address range. The priority should be identical on each node and be a higher value than any other configured RP candidate that is not a member of this rp-set. Although there is no set maximum of addresses that can be configured in an rp-set, up to 15 multicast addresses is recommended. The no form of the command removes an entry from the list.
Parameters	<i>ipv6-address</i> — Specifies the address used by the other node as the RP candidacy address for the same multicast group address range as configured on this node. Values <i>ipv6-address</i> : x:x:x:x:x:x:x (eight 16-bit pieces)

x::x::x::x::d.d.d
 x [0 to FFFF]H
 d [0 to 255]D

embedded-rp

Syntax	embedded-rp
Context	config>service>vprn>pim>rp>ipv6
Description	This command enables context to configure IPv6 embedded RP parameters.

group-range

Syntax	[no] group-range {ipv6-address/prefix-length}
Context	config>service>vprn>pim>rp>ipv6>embedded-rp config>service>vprn>pim>rp>ipv6>rp-candidate
Description	This command configures the group address or range of group addresses for which this router can be the rendezvous point (RP). Use the no form of this command to remove the group address or range of group addresses for which this router can be the RP from the configuration.
Parameters	<i>ipv6-address</i> — Specifies the addresses or address ranges that this router can be an RP. <i>prefix-length</i> — Specifies the address prefix length.

Values

ipv6-address	: x::x::x::x::x (eight 16-bit pieces)
	x::x::x::x::d.d.d
	x [0 to FFFF]H
	d [0 to 255]D
prefix-length	[8 to 128] // for embedded-rp
prefix-length	[16 to 128] // for rp-candidate

group-prefix

Syntax	[no] group-prefix grp-ipv6-address/prefix-length
Context	config>service>vprn>pim>rp>ipv6>static
Description	The group-prefix for a static-rp defines a range of multicast-ip-addresses for which this static RP is applicable.

The **no** form of the command removes the criterion.

Parameters *grp-ipv6-address* — Specifies the multicast IPv6 address.
prefix-length — Specifies the address prefix length.

Values

grp-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	[8 to 128]

rp-candidate

Syntax **rp-candidate**

Context config>service>vprn>pim>rp
config>service>vprn>pim>rp>ipv6

Description This command enters the context to configure the candidate rendezvous point (RP) parameters.

Default enabled when PIM is enabled

group-range

Syntax [**no**] **group-range** {*ip-prefix/mask* | *ip-prefix netmask*}

Context config>service>vprn>pim>rp>rp-candidate
config>service>vprn>pim>ssm

Description This command configures the group address or range of group addresses for which this router can be the rendezvous point (RP).

Use the **no** form of this command to remove the group address or range of group addresses for which this router can be the RP from the configuration.

Parameters *ip-prefix* — Specifies the addresses or address ranges that this router can be an RP.

Values

ipv4-prefix	- a.b.c.d
ipv4-prefix-le	- [0 to 32]
ipv6-prefix	- x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d
x	- [0 to FFFF]H
d	- [0 to 255]D
ipv6-prefix-le	- [0 to 128]

mask — Specifies the address mask with the address to define a range of addresses.

netmask — Specifies the subnet mask in dotted decimal notation.

Values :a.b.c.d (network bits all 1 and host bits all 0)

holdtime

Syntax	holdtime <i>holdtime</i> no holdtime
Context	config>service>vprn>pim>rp>rp-candidate config>service>vprn>pim>rp>ipv6>rp-candidate
Description	Use this command to define the length of time neighboring router consider this router to be up. Use the no form of this command to revert to the default value.
Default	holdtime 150
Parameters	<i>holdtime</i> — Specifies the length of time, in seconds, that neighbor should consider the sending router to be operational. Values 0 to 255

priority

Syntax	priority <i>priority</i> no priority
Context	config>service>vprn>pim>rp>rp-candidate
Description	This command defines the priority used to become the rendezvous point (RP). The higher the priority value, the more likely that this router will become the RP. Use the no form of this command to revert to the default value.
Default	priority 192
Parameters	<i>priority</i> — Specifies the priority to become the designated router. The higher the value the more likely the router will become the RP. Values 0 to 255

static

Syntax	static
Context	config>service>vprn>pim>rp

Description This command enables access to the context to configure a static rendezvous point (RP) of a PIM-SM protocol instance.

address

Syntax `[no] address ip-address`

Context `config>service>vprn>pim>rp>static`

Description This command configures the static rendezvous point (RP) address.

The **no** form of this command removes the static RP entry from the configuration.

group-prefix

Syntax `[no] group-prefix {grp-ip-address/mask | grp-ip-address netmask}`

Context `config>service>vprn>pim>rp>static`

Description The **group-prefix** for a static-rp defines a range of multicast-ip-addresses for which a certain RP is applicable.

The **no** form of the command removes the criterion.

Parameters *grp-ip-address* — Specifies the multicast IP address.

mask — Defines the mask of the multicast-ip-address.

Values 4 to 32

netmask — The subnet mask in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

override

Syntax `[no] override`

Context `config>service>vprn>pim>rp>static`

Description This command changes the precedence of static RP over dynamically learned Rendezvous Point (RP).

When enabled, the static group-to-RP mappings take precedence over the dynamically learned mappings.

Default no override

rpf-table

Syntax	rpf-table { rtable-m rtable-u both } no rpf-table
Context	config>service>vprn>pim config>router>msdp
Description	This command configures the sequence of route tables used to find a Reverse Path Forwarding (RPF) interface for a particular multicast route. By default, only the unicast route table is looked up to calculate RPF interface towards the source/rendezvous point. However, the operator can specify the following: • use the unicast route table only • use the multicast route table only • use both the route tables
Default	rpf-table rtable-u
Parameters	rtable-m — Specifies that only the multicast route table will be used by the multicast protocol (PIM) for IPv4 RPF checks. This route table will contain routes submitted by static routes, IS-IS and OSPF. rtable-u — Specifies only that the unicast route table will be used by the multicast protocol (PIM) for IPv4 RPF checks. This route table will contain routes submitted by all the unicast routing protocols. both — Specifies that the multicast route table will be used first by the multicast protocol (PIM) for checks, and then the unicast route table will be used if the multicast route table lookup fails. rtable-m is checked before rtable-u.

spt-switchover-threshold

Syntax	spt-switchover-threshold { <i>grp-ip-address/mask</i> <i>grp-ip-address netmask</i> } <i>spt-threshold</i> no spt-switchover-threshold { <i>grp-ip-address/mask</i> <i>grp-ip-address netmask</i> }
Context	config>service>vprn>pim
Description	This command configures a shortest path tree (SPT tree) switchover threshold for a group prefix.
Parameters	<i>grp-ip-address</i> — Specifies the multicast group address. <i>mask</i> — Defines the mask of the multicast-ip-address. Values 4 to 32 <i>netmask</i> — The subnet mask in dotted decimal notation. Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

spt-threshold — Specifies the configured threshold in kilobits per second (kb/s) for the group to which this (S,G) belongs. For a group G configured with a threshold, switchover to SPT for an (S,G) is attempted only if the (S,G)'s rate exceeds this configured threshold.

ssm-assert-compatible-mode

Syntax	ssm-assert-compatible-mode [enable disable]
Context	config>service>vprn>pim
Description	This command specifies whether SSM assert is enabled in compatibility mode for this PIM protocol instance. When enabled, for SSM groups, PIM will consider the SPT bit to be implicitly set to compute the value of CouldAssert (S,G,I) as defined in RFC 4601, <i>Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)</i> . When disabled, for SSM groups, PIM will not assume the SPT bit to be set. The SPT bit will be set by Update_SPTbit(S,G,iif) macro defined in RFC 4601.
Default	ssm-assert-compatible-mode disable
Parameters	enable — enables SSM assert in compatibility mode for this PIM protocol instance disable — disabled SSM assert in compatibility mode for this PIM protocol instance

ssm-default-range-disable

Syntax	ssm-default-range-disable ipv4
Context	config>service>vprn>pim
Description	This command specifies whether to disable the use of default range (232/8) for SSM so that it can be used by ASM to process (*,G). When enabled, the use of default range is disabled for SSM and it can be used by ASM. When disabled, the SSM default range is enabled.
Default	ssm-default-range-disable

ssm-groups

Syntax	[no] ssm-groups
Context	config>service>vprn
Description	This command enables access to the context to enable a source-specific multicast (SSM) configuration instance.

3.8.2.33 PPPoE Commands



Note: The commands described in this section apply only to the 7450 ESS and 7750 SR.

pppoe

Syntax	[no] pppoe
Context	config>service>vprn>sub-if>grp-if
Description	This command enters the context to configure PPPoE parameters.

dhcp-client

Syntax	dhcp-client
Context	config>service>vprn>sub-if>grp-if>pppoe
Description	This command enters the context to configure the PPPoE-to-DHCP options.

ccag-use-origin-sap

Syntax	[no] ccag-use-origin-sap
Context	config>service>vprn>sub-if>grp-if>pppoe>dhcp-client
Description	This command enables the original VPLS SAP to be included in the circuit-id option to send to the DHCP server (in case this interface is connected to a VPLS by a CCA MDA). The no form of the command disables the feature.
Default	no ccag-use-origin-sap

pap-chap-user-db

Syntax	pap-chap-user-db <i>local-user-db-name</i> no pap-chap-user-db
Context	config>service>vprn>sub-if>grp-if>pppoe
Description	This command configures the local user database to use for PPP Challenge-Handshake Authentication Protocol/Password Authentication Protocol (PAP/CHAP) authentication.

If an authentication policy is also configured, **pppoe-access-method** must be set to none in this authentication policy to use the local user database (in that case RADIUS authentication will not be used for PPPoE hosts).

Parameters *local-user-db-name* — Specifies the local user database to use for authentication.

policy

Syntax **policy** *pppoe-policy-name*
no policy

Context config>service>vprn>sub-if>grp-if>pppoe

Description This command associates a PPPoE policy on this interface.

Default policy default

Parameters *pppoe-policy-name* — Specifies a PPPoE policy up to 32 characters in length on this interface.

sap-session-limit

Syntax **sap-session-limit** *sap-session-limit*
no sap-session-limit

Context config>service>vprn>sub-if>grp-if>pppoe

Description This command specifies the number of PPPoE hosts per SAP allowed for this group-interface.

Default sap-session-limit 1

Parameters *sap-session-limit* — Specifies the number of PPPoE hosts per SAP allowed.

Values 1 to 20000

session-limit

Syntax **session-limit** *session-limit*
no session-limit

Context config>service>vprn>sub-if>grp-if>pppoe

Description This command specifies the number of PPPoE hosts allowed for this group interface.

Default session-limit 1

Parameters	<i>session-limit</i> — Specifies the number of PPPoE hosts allowed.
Values	1 to 20000

3.8.2.34 MSDP Configuration Commands

msdp

Syntax	[no] msdp
Context	config>service>vprn
Description	This command enables a Multicast Source Discovery Protocol (MSDP) instance. When an MSDP instance is created, the protocol is enabled. To start or suspend execution of the MSDP protocol without affecting the configuration, use the [no] shutdown command. For the MSDP protocol to function, at least one peer must be configured. When MSDP is configured and started, an appropriate event message should be generated. When the no form of the command is executed, all sessions must be terminated and an appropriate event message should be generated. When all peering sessions are terminated, an event message per peer is not required. The no form of the command deletes the MSDP protocol instance, removing all associated configuration parameters.
Default	no msdp

active-source-limit

Syntax	active-source-limit <i>number</i> no active-source-limit
Context	config>service>vprn>msdp config>service>vprn>msdp>group config>service>vprn>msdp>group>peer config>service>vprn>msdp>peer config>service>vprn>msdp>source
Description	This option controls the maximum number of active source messages that will be accepted by Multicast Source Discovery Protocol (MSDP), effectively controlling the number of active sources that can be stored on the system.

The **no** form of this command reverts the number of source message limit to default operation.

Default	no active-source-limit
Parameters	<i>number</i> — Defines how many active sources can be maintained by MSDP.
Values	0 to 1000000

data-encapsulation

Syntax	[no] data-encapsulation
Context	config>service>vprn>msdp
Description	This command configures a rendezvous point (RP) using Multicast Source Discovery Protocol (MSDP) to encapsulate multicast data received in MSDP register messages inside forwarded MSDP source-active messages.
Default	data-encapsulation

export

Syntax	export <i>policy-name</i> [<i>policy-name...</i>(up to 5 max)] no export
Context	config>service>vprn>msdp config>service>vprn>msdp>group config>service>vprn>msdp>group>peer config>service>vprn>msdp>peer
Description	This command specifies the policies to export source active state from the source active list into Multicast Source Discovery Protocol (MSDP). If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified. The no form of the command removes all policies from the configuration.
Default	no export
Parameters	<i>policy-name</i> — Specifies the export policy name, up to 32 characters. Up to five policy-name arguments can be specified.

If you configure an export policy at the global level, each individual peer inherits the global policy. If you configure an export policy at the group level, each individual peer in a group inherits the group's policy. If you configure an export policy at the peer level, then policy only applies to the peer where it is configured.

group

Syntax	[no] group <i>group-name</i>
Context	config>service>vprn>msdp
Description	This command enables access to the context to create or modify a Multicast Source Discovery Protocol (MSDP) group. To configure multiple MSDP groups, include multiple group statements. By default, the group's options are inherited from the global MSDP options. To override these global options, group-specific options within the group statement can be configured. If the group name provided is already configured then this command only provides the context to configure the options pertaining to this group. If the group name provided is not already configured, then the group name must be created and the context to configure the parameters pertaining to the group should be provided. In this case, the \$ prompt to indicate that a new entity (group) is being created should be used. For a group to be of use, at least one peer must be configured.
Default	no group
Parameters	<i>group-name</i> — Specifies a unique name for the MSDP group.

import

Syntax	import <i>policy-name</i> [<i>policy-name...</i> (up to 5 max)] no import
Context	config>service>vprn>msdp config>service>vprn>msdp>group config>service>vprn>msdp>group>peer config>service>vprn>msdp>peer
Description	This command specifies the policies to import source active state from Multicast Source Discovery Protocol (MSDP) into source active list. 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 import commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

If you configure an import policy at the global level, each individual peer inherits the global policy.

If you configure an import policy at the group level, each individual peer in a group inherits the group's policy.

If you configure an import policy at the peer level, then policy only applies to the peer where it is configured.

The **no** form of the command removes all policies from the configuration.

Default no import

Parameters *policy-name* — Specifies the import policy name. Up to five policy-name arguments can be specified.

local-address

Syntax **local-address** *address*
no local-address

Context config>service>vprn>msdp
config>service>vprn>msdp>group
config>service>vprn>msdp>group>peer
config>service>vprn>msdp>peer

Description This command configures the local end of a Multicast Source Discovery Protocol (MSDP) session. For MSDP to function, at least one peer must be configured. When configuring a peer, you must include this local-address command to configure the local end of the MSDP session. This address must be present on the node and is used to validate incoming connections to the peer and to establish connections to the remote peer.

If the user enters this command, then the address provided is validated and will be used as the local address for MSDP peers from that point. If a subsequent local-address command is entered, it will replace the existing configuration and existing sessions will be terminated.

Similarly, when the **no** form of this command is entered, the existing local-address will be removed from the configuration and the existing sessions will be terminated.

Whenever a session is terminated, all information pertaining to and learned from that peer will be removed.

Whenever a new peering session is created or a peering session is lost, an event message should be generated.

The **no** form of this command removes the local-address from the configuration.

Default no local-address

Parameters *address* — Specifies an existing address on the node.

mode

Syntax	mode { mesh-group standard }
Context	config>service>vprn>msdp>group
Description	This command configures groups of peers in a full mesh topology to limit excessive flooding of source-active messages to neighboring peers. Multicast Source Discovery Protocol (MSDP) peers can be configured grouped in a full-mesh topology that prevents excessive flooding of source-active messages to neighboring peers. In a meshed configuration, all members of the group must have a peer connection with every other mesh group member. If this rule is not adhered to, then unpredictable results may occur.
Default	mode standard
Parameters	mesh-group — Specifies that source-active message received from a mesh group member are always accepted but are not flooded to other members of the same mesh group. These source-active messages are only flooded to non-mesh group peers or members of other mesh groups. standard — Specifies a non-meshed mode.

peer

Syntax	[no] peer <i>peer-address</i>
Context	config>service>vprn>msdp config>service>vprn>msdp>group
Description	This command configures peer parameters. Multicast Source Discovery Protocol (MSDP) must have at least one peer configured. A peer is defined by configuring a local-address that can be used by this node to set up a peering session and the address of a remote MSDP router. It is the address of this remote peer that is configured in this command and it identifies the remote MSDP router address. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. It may be required to have multiple peering sessions in which case multiple peer statements should be included in the configurations. By default, the options applied to a peer are inherited from the global or group-level. To override these inherited options, include peer-specific options within the peer statement. If the peer address provided is already a configured peer, then this command only provides the context to configure the parameters pertaining to this peer.

If the peer address provided is not already a configured peer, then the peer instance must be created and the context to configure the parameters pertaining to this peer should be provided. In this case, the \$ prompt to indicate that a new entity (peer) is being created should be used.

The peer address provided will be validated and, if valid, will be used as the remote address for an MSDP peering session.

When the **no** form of this command is entered, the existing peering address will be removed from the configuration and the existing session will be terminated. Whenever a session is terminated, all source active information pertaining to and learned from that peer will be removed. Whenever a new peering session is created or a peering session is lost, an event message should be generated.

At least one peer must be configured for MSDP to function.

Parameters *peer-address* — The address configured in this statement must identify the remote MSDP router that the peering session must be established with.

receive-msdp-msg-rate

Syntax	receive-msdp-msg-rate <i>number interval seconds</i> [threshold <i>number</i>] no receive-msdp-msg-rate
Context	config>service>vprn>msdp config>service>vprn>msdp>group config>service>vprn>msdp>group>peer config>service>vprn>msdp>peer
Description	This command limits the number of Multicast Source Discovery Protocol (MSDP) messages that are read from the TCP session. It is possible that an MSDP/ RP router may receive a large number of MSDP protocol message packets in a particular source active message. After the number of MSDP packets (including source active messages) defined in the threshold have been processed, the rate of all other MSDP packets is rate limited by no longer accepting messages from the TCP session until the time (seconds) has elapsed. The no form of this command reverts this active-source limit to default operation.
Default	no receive-msdp-msg-rate
Parameters	<i>number</i> — Defines the number of MSDP messages (including source active messages) that are read from the TCP session per the number of seconds. Values 10 to 10000 Default 0

interval seconds — Defines the time that, together with the *number* parameter, defines the number of MSDP messages (including source active messages) that are read from the TCP session within the configured number of seconds.

Values 1 to 600

Default 0

threshold number — The number of MSDP messages can be processed before the MSDP message rate limiting function described above is activated; this is particularly of use during at system startup and initialization.

Values 1 to 1000000

Default 0

rpf-table

Syntax	rpf-table {rtable-m rtable-u both} no rpf-table
Context	config>service>vprn>msdp
Description	This command configures the sequence of route tables used to find a Reverse Path Forwarding (RPF) interface for a particular multicast route. By default, only the unicast route table is looked up to calculate RPF interface towards the source/rendezvous point. However, the operator can specify the following: • use the unicast route table only • use the multicast route table only or • use both the route tables The no form of this command reverts to the default.
Default	rpf-table rtable-u
Parameters	rtable6-m — Specifies that only the multicast route table will be used by the multicast protocol (PIM) for IPv4 RPF checks. This route table will contain routes submitted by static routes, ISIS and OSPF. rtable6-u — Specifies only that the unicast route table will be used by the multicast protocol (PIM) for IPv4 RPF checks. This route table will contain routes submitted by all the unicast routing protocols. both — Will always look up first in the multicast route table and, if there is a route, it will use it. If PIM does not find a route in the first lookup, it will try to find it in the unicast route table. Rtable-m is checked before rtable6-u.

sa-timeout

Syntax	sa-timeout <i>seconds</i> no sa-timeout
Context	config>service>vprn>msdp
Description	This command configures the value for the SA entries in the cache. If these entries are not refreshed within the timeout value, they are removed from the cache. Normally, the entries are refreshed at least once a minute. But under high load with many of MSDP peers, the refresh cycle could be incomplete. A higher timeout value (more than 90) could be useful to prevent instabilities in the MSDP cache.
Default	90
Parameters	<i>seconds</i> — Specifies the time, in seconds, to wait for a response from the peer before declaring the peer unavailable. Values 90 to 600

source

Syntax	[no] source <i>ip-prefix/mask</i>
Context	config>service>vprn>msdp
Description	This command limits the number of active source messages the router accepts from sources in the specified address range. If the prefix and mask provided is already a configured then this command only provides the context to configure the parameters pertaining to this active source-message filter. If the prefix and mask provided is not already a configured, then the source node instance must be created and the context to configure the parameters pertaining to this node should be provided. In this case, the \$ prompt to indicate that a new entity (source) is being created should be used. The source active msdp messages are not rate limited based on the source address range. The no form of this message removes the source active rate limiter for this source address range.
Parameters	<i>ip-prefix</i> — The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area. Values ip-prefix/mask: ip-prefix a.b.c.d (host bits must be 0) <i>mask</i> — The subnet mask for the range expressed as a decimal integer mask length or in dotted decimal notation. Values 0 to 32 (mask length), 0.0.0.0 to 255.255.255.255 (dotted decimal)

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>service>vprn>msdp>group>peer config>service>vprn>msdp>peer
Description	This command configures a Message Digest 5 (MD5) authentication key to be used with a specific Multicast Source Discovery Protocol (MSDP) peering session. The authentication key must be configured per peer as such no global or group configuration is possible.
Default	no authentication-key (All MSDP messages are accepted and the MD5 signature option authentication key is disabled.)
Parameters	<i>authentication-key</i> — The authentication key. Allowed values are any string up to 256 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), enclose the entire string in quotation marks (" "). <i>hash-key</i> — The hash key. The key can be any combination of ASCII characters up to 451 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" "). This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided. hash — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified. hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the hash2 encrypted variable cannot be copied and pasted. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.

default-peer

Syntax	default-peer no default-peer
Context	config>service>vprn>msdp>group>peer config>service>vprn>msdp>peer
Description	Using the default peer mechanism, a peer can be selected as the default Multicast Source Discovery Protocol (MSDP) peer. As a result, all source-active messages from the peer will be accepted without the usual peer-reverse-path-forwarding (RPF) check.

The MSDP peer-RPF check is different from the normal multicast RPF checks. The peer-RPF check is used to stop source-active messages from looping. A router validates source-active messages originated from other routers in a deterministic fashion.

A set of rules is applied in order to validate received source-active messages, and the first rule that applies determines the peer-RPF neighbor. All source-active messages from other routers are rejected. The rules applied to source-active messages originating at Router S received at Router R from Router N are as follows:

- If Router N and router S are one and the same, then the message is originated by a direct peer-RPF neighbor and will be accepted.
- If Router N is a configured peer, or a member of the Router R mesh group then its source-active messages are accepted.
- If Router N is the Border Gateway Protocol (BGP) next hop of the active multicast RPF route toward Router S then Router N is the peer-RPF neighbor and its source-active messages are accepted.
- If Router N is an external BGP peer of Router R and the last autonomous system (AS) number in the BGP AS-path to Router S is the same as Router N's AS number, then Router N is the peer-RPF neighbor, and its source-active messages are accepted.
- If Router N uses the same next hop as the next hop to Router S, then Router N is the peer-RPF neighbor, and its source-active messages are accepted.
- If Router N fits none of the above rules, then Router N is not a peer-RPF neighbor, and its source-active messages are rejected.

Default no default-peer

3.8.2.35 MLD Configuration Commands

mld

Syntax	[no] mld
Context	config>service>vprn
Description	This command enters the context to configure Multicast Listener Discovery (MLD) parameters. The no form of the command disables MLD.
Default	no mld

interface

Syntax	[no] interface <i>ip-int-name</i>

Context	config>service>vprn>mld
Description	This command enters the context to configure an Multicast Listener Discovery (MLD) interface. The interface is a local identifier of the network interface on which reception of the specified multicast address is to be enabled or disabled. The no form of the command deletes the MLD interface. The shutdown command in the config>router>mld>interface context can be used to disable an interface without removing the configuration for the interface.
Default	no interface
Parameters	<i>ip-int-name</i> — The IP interface name. 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 any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. If the IP interface name does not exist or does not have an IP address configured an error message will be returned. If the IP interface exists in a different area it will be moved to this area.

disable-router-alert-check

Syntax	[no] disable-router-alert-check
Context	config>service>vprn>mld>interface
Description	This command disables router alert checking for MLD messages received on this interface. The no form of the command enables the router alert checking.

import

Syntax	import <i>policy-name</i> no import
Context	config>service>vprn>mld>interface
Description	This command specifies the import route policy to be used for determining which membership reports are accepted by the router. Route policies are configured in the config>router>policy-options context. When an import policy is not specified, all the MLD reports are accepted. The no form of the command removes the policy association from the MLD instance.
Default	no import

Parameters *policy-name* — The route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. Route policies are configured in the **config>router>policy-options** context.

max-groups

Syntax **max-groups** *value*
no max-groups

Context config>service>vprn>mld>interface

Description This command specifies the maximum number of groups for which MLD can have local receiver information based on received MLD reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed.

Default 0 (no limit to the number of groups)

Parameters *value* — Specifies the maximum number of groups for this interface.

Values 1 to 16000

query-interval

Syntax **query-interval** *seconds*
no query-interval

Context config>service>vprn>mld>interface

Description This command specifies the frequency that the querier router transmits general host-query messages. The host-query messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.

Default 125

Parameters *seconds* — The time frequency, in seconds, that the router transmits general host-query messages.

Values 2 to 1024

query-last-member-interval

Syntax **query-last-member-interval** *seconds*

Context config>service>vprn>mld>interface

Description	This command configures the frequency at which the querier sends group-specific query messages including messages sent in response to leave-group messages. The lower the interval, the faster the detection of the loss of the last member of a group.
Default	1
Parameters	<i>seconds</i> — Specifies the frequency, in seconds, at which query messages are sent. Values 1 to 1024

query-response-interval

Syntax	query-response-interval <i>seconds</i>
Context	config>service>vprn>mld>interface
Description	This command specifies how long the querier router waits to receive a response to a host-query message from a host.
Default	10
Parameters	<i>seconds</i> — Specifies the length of time to wait to receive a response to the host-query message from the host. Values 1 to 1023

static

Syntax	static
Context	config>service>vprn>mld>interface
Description	This command tests multicast forwarding on an interface without a receiver host. When enabled, data is forwarded to an interface without receiving membership reports from host members.

group

Syntax	[no] group <i>grp-ipv6-address</i> [no] group start <i>grp-ipv6-address</i> end <i>grp-ipv6-address</i> [step <i>ipv6-address</i>]
Context	config>service>vprn>mld>interface>static
Description	This command enters the context to add a static multicast group either as a (*,G) or one or more (S,G) records. Use MLD static group memberships to test multicast forwarding without a receiver host. When MLD static groups are enabled, data is forwarded to an interface without receiving membership reports from host members.

When static MLD group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static MLD group entries do not generate join messages toward the RP.

The **no** form of the command removes the IPv6 address from the configuration.

Parameters *grp-ipv6-address* — Specifies an MLD multicast group address that receives data on an interface. The IP address must be unique for each static group.

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

start *grp-ipv6-address* — Specifies the start multicast group 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

end *grp-ipv6-address* — Specifies the end multicast group 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

step *ipv6-address* — Specifies the step increment.

SOURCE

Syntax [**no**] **source** *src-ipv6-address*

Context config>service>vprn>mld>interface>static>group

Description This command specifies an IPv6 unicast address that sends data on an interface. This enables a multicast receiver host to signal a router the group to receive multicast traffic from, and from the sources that the traffic is expected.

The **source** command is mutually exclusive with the specification of individual sources for the same group.

The **source** command, in combination with the group, is used to create a specific (S,G) static group entry.

The **no** form of the command removes the source from the configuration.

Parameters *src-ipv6-address* — Specifies the IPv6 unicast address.

starg

Syntax [no] **starg**

Context config>service>vprn>mld>interface>static>group

Description This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified.

Use the **no** form of the command to remove the **starg** entry from the configuration.

version

Syntax **version** *version*
 no version

Context config>service>vprn>mld>interface

Description This command specifies the MLD version. If routers run different versions, they will negotiate the lowest common version of MLD that is supported by hosts on their subnet and operate in that version. For MLD to function correctly, all routers on a LAN should be configured to run the same version of MLD on that LAN.

Default version 2

Parameters *version* — Specifies the MLD version number.

Values 1, 2

robust-count

Syntax **robust-count** *robust-count*
 no robust-count

Context config>service>vprn>mld

Description This command configures the robust count. The robust-count variable allows tuning for the expected packet loss on a subnet. If a subnet anticipates losses, the robust-count variable can be increased.

Default robust-count 2

Parameters *robust-count* — Specifies the robust count value.

Values 2 to 10

ssm-translate

Syntax **ssm-translate**

Context config>service>vprn>mld

Description This command enters the context to configure group ranges which are translated to SSM (S,G) entries. If the static entry needs to be created, it has to be translated from a IGMPv1 IGMPv2 request to a Source Specific Multicast (SSM) join. An SSM translate source can only be added if the **starg** command is not enabled. An error message is generated if you try to configure the **source** command with **starg** command enabled.

grp-range

Syntax [**no**] **grp-range** *start end*

Context config>service>vprn>mld>ssm-translate

Description This command is used to configure group ranges which are translated to SSM (S,G) entries.

Parameters *start* — An IP address that specifies the start of the group range.
end — An IP address that specifies the end of the group range. This value should always be greater than or equal to the value of the *start* value.

source

Syntax [**no**] **source** *ip-address*

Context config>service>vprn>mld>ssm-translate>grp-range

Description This command specifies the source IP address for the group range. Whenever a (*,G) report is received in the range specified by **grp-range** *start* and *end* parameters, it is translated to an (S,G) report with the value of this object as the source address.

Parameters *ip-address* — Specifies the IP address that will be sending data.

3.8.2.36 RIP Commands

rip

Syntax	[no] rip
Context	config>service>vprn
Description	This command enables the RIP protocol on the given VPRN IP interface. The no form of the command disables the RIP protocol from the given VPRN IP interface.
Default	no rip

ripng

Syntax	[no] ripng
Context	config>router
Description	This command creates the context to configure the RIPng protocol instance. When a RIPng instance is created, the protocol is enabled by default. To start or suspend execution of the RIP protocol without affecting the configuration, use the [no] shutdown command. The no form of the command deletes the RIP protocol instance removing all associated configuration parameters.
Default	no ripng

authentication-key

Syntax	authentication-key <i>[authentication-key hash-key]</i> [hash hash2] no authentication-key
Context	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor
Description	This command sets the authentication password to be passed between RIP neighbors. The authentication type and authentication key must match exactly for the RIP message to be considered authentic and processed. The no form of the command removes the authentication password from the configuration and disables authentication.

Default	no authentication-key
Parameters	<i>authentication-key</i> — The authentication key. The key can be any combination of ASCII characters up to 16 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" "). <i>hash-key</i> — The hash key. The key can be any combination of ASCII characters up to 33 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" "). This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided. hash — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the hash2 encrypted variable cannot be copied and pasted. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.

authentication-type

Syntax	authentication-type {none password message-digest} no authentication-type
Context	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor
Description	This command defines the type of authentication to be used between RIP neighbors. The type and password must match exactly for the RIP message to be considered authentic and processed. The no form of the command removes the authentication type from the configuration and effectively disables authentication.
Default	no authentication-type
Parameters	<i>none</i> — No authentication is used. <i>simple</i> — A simple clear-text password is sent. <i>md5</i> — MD5 authentication is used.

bfd-enable

Syntax	[no] bfd-enable
Context	config> service>vprn>rip>group>neighbor config> service>vprn>ripng>group>neighbor
Description	This command enables bi-directional forwarding (BFD) to control the state of the associated protocol adjacency. By enabling BFD on a given protocol interface, the state of the RIP neighbor 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 configuration context. The no form of this command removes BFD from the associated IGP/BGP protocol adjacency.

check-zero

Syntax	check-zero {enable disable} no check-zero
Context	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor config>service>vprn>ripng config>service>vprn>ripng>group config>service>vprn>ripng>group>neighbor
Description	This command enables checking for zero values in fields specified to be zero by the RIPv1 and RIPv2 specifications. The no form of the command disables this check and allows the receipt of RIP messages even if the mandatory zero fields are non-zero.
Default	no check-zero
Parameters	enable — Enables checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications and rejecting non-compliant RIP messages. disable — Disables the checking and allows the receipt of RIP messages even if the mandatory zero fields are non-zero.

split-horizon

Syntax	split-horizon {enable disable} no split-horizon
Context	config>service>vprn>rip

```
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor
config>service>vprn>ripng
config>service>vprn>ripng>group
config>service>vprn>ripng>group>neighbor
```

Description This command enables the use of split-horizon. RIP uses split-horizon with poison-reverse to protect from such problems as “counting to infinity”. Split-horizon with poison reverse means that routes learned from a neighbor through a given interface are advertised in updates out of the same interface but with a metric of 16 (infinity).

The **split-horizon disable** command enables split horizon without poison reverse. This allows the routes to be re-advertised on interfaces other than the interface that learned the route, with the advertised metric equaling an increment of the metric-in value.

This configuration parameter can be set at three levels: global level (applies to all groups and neighbor interfaces), group level (applies to all neighbor interfaces in the group) or neighbor level (only applies to the specified neighbor interface). The most specific value is used. In particular if no value is set (**no split-horizon**), the setting from the less specific level is inherited by the lower level.

The **no** form of the command disables split horizon command which allows the lower level to inherit the setting from an upper level.

Default enabled

export

Syntax **export** *policy-name* [*policy-name...*(up to 5 max)]
no export

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor
config>service>vprn>ripng
config>service>vprn>ripng>group
config>service>vprn>ripng>group>neighbor

Description This command specifies the export route policies used to determine which routes are exported to RIP. If no export policy is specified, non-RIP routes will not be exported from the routing table manager to RIP; RIP-learned routes will be exported to RIP neighbors.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of the command removes all policies from the configuration.

Default no export

Parameters *policy-name* — The export route policy name. Allowed values are any string up to 32 characters in length and composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the string must be enclosed between double quotes. The specified names must already be defined.

export-limit

Syntax **export-limit** *number* [**log** *percentage*]
no export-limit

Context config>service>vprn>rip
 config>service>vprn>ripng

Description This command configures the maximum number of routes (prefixes) that can be exported into RIP from the route table.

 The **no** form of the command removes the parameters from the configuration.

Default no export-limit

Parameters *number* — Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.

Values 1 to 4294967295

log *percentage* — Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.

Values 1 to 100

import

Syntax **import** *policy-name* [*policy-name...*(up to 5 max)]
no import

Context config>service>vprn>rip
 config>service>vprn>rip>group
 config>service>vprn>rip>group>neighbor
 config>service>vprn>ripng
 config>service>vprn>ripng>group
 config>service>vprn>ripng>group>neighbor

Description This command configures import route policies to determine routes that will be accepted from RIP neighbors. If no import policy is specified, RIP accepts all routes from configured RIP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

If multiple policy names are specified, the policies are evaluated in the order that they are specified. The first policy that matches is applied. If multiple import commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of the command removes all policies from the configuration.

Default	no import
Parameters	<i>policy-name</i> — The import route policy name. Allowed values are any string up to 32 characters in length and composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes. The specified names must already be defined.

message-size

Syntax	message-size <i>max-num-of-routes</i> no message-size				
Context	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor config>service>vprn>ripng config>service>vprn>ripng>group config>service>vprn>ripng>group>neighbor				
Description	This command sets the maximum number of routes per RIP update message. The no form of the command resets the maximum number of routes back to the default of 25.				
Default	no message-size				
Parameters	<i>size</i> — An Integer. <table> <tr> <td>Values</td><td>25 to 255</td></tr> <tr> <td>Default</td><td>25</td></tr> </table>	Values	25 to 255	Default	25
Values	25 to 255				
Default	25				

metric-in

Syntax	metric-in <i>metric</i> no metric-in
Context	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor config>service>vprn>ripng config>service>vprn>ripng>group config>service>vprn>ripng>group>neighbor

Description	This command sets the metric added to routes that were received from a RIP neighbor. The no form of the command reverts the <i>metric</i> value back to the default.
Default	no metric-in
Parameters	<i>metric</i> — The value added to the metric of routes received from a RIP neighbor, expressed as a decimal integer. Values 1 to 16

metric-out

Syntax	metric-out <i>metric</i> no metric-out
Context	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor config>service>vprn>ripng config>service>vprn>ripng>group config>service>vprn>ripng>group>neighbor
Description	This command sets the metric added to routes that were exported into RIP and advertised to RIP neighbors. The no form of the command removes the command from the config and resets the metric-in value back to the default.
Default	no metric-out
Parameters	<i>metric</i> — The value added to the metric for routes exported into RIP and advertised to RIP neighbors, expressed as a decimal integer. Values 1 to 16

preference

Syntax	preference <i>preference</i> no preference
Context	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor config>service>vprn>ripng config>service>vprn>ripng>group config>service>vprn>ripng>group>neighbor

Description This command sets the route preference assigned to RIP routes. This value can be overridden by route policies.

The **no** form of the command resets the *preference* to the default.

Default no preference

Parameters *preference* — An integer.

Values 1 to 255

Default 100

propagate-metric

Syntax [**no**] **propagate-metric**

Context config>service>vprn>rip
config>service>vprn>ripng

Description This command enables the BGP MED to be used to configure the RIP metric at the BGP to RIP transition on egress routers. BGP always configures the BGP MED to the RIP metric at the ingress router. When **propagate-metric** is configured, the RIP metric at egress routers is configured as the BGP MED attribute added to the optional value configured with the **metric-out** command.

The **no** version of this command sets the RIP metric to the optional value configured with the metric-out command plus 1.

Default no propagate-metric

receive

Syntax **receive {both | none | version-1 | version-2}**
no receive

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor
config>service>vprn>ripng
config>service>vprn>ripng>group
config>service>vprn>ripng>group>neighbor

Description This command configures the type(s) of RIP updates that will be accepted and processed.

If **both** or **version-2** is specified, the RIP instance listens for and accepts packets sent to the broadcast and multicast (224.0.0.9) addresses.

If **version-1** is specified, the router only listens for and accepts packets sent to the broadcast address.

This control can be issued at the global, group or interface level. The default behavior accepts and processes both RIPv1 and RIPv2 messages.

The **no** form of the command resets the type of messages accepted to both.

Default no receive

Parameters *both* — Receive RIP updates in either Version 1 or Version 2 format.
none — Do not accept and RIP updates.
version-1 — Router should only accept RIP updates in Version 1 format.
version-2 — Router should only accept RIP updates in Version 2 format.

send

Syntax **send {broadcast | multicast | none | version-1 | both}**
no send

Context config>service>vprn>rip
 config>service>vprn>rip>group
 config>service>vprn>rip>group>neighbor
 config>service>vprn>ripng
 config>service>vprn>ripng>group
 config>service>vprn>ripng>group>neighbor

Description This command specifies the type of RIP messages sent to RIP neighbors. This control can be issued at the global, group or interface level. The default behavior sends RIPv2 messages with the multicast (224.0.0.9) destination address.

If **version-1** is specified, the router only listens for and accepts packets sent to the broadcast address.

The **no** form of this command resets the type of messages sent back to the default value.

Default no send

Parameters *broadcast* — Send RIPv2 formatted messages to the broadcast address.
multicast — Send RIPv2 formatted messages to the multicast address.
none — Do not send any RIP messages (i.e. silent listener).
version-1 — Send RIPv1 formatted messages to the broadcast address.
both — Send both RIP v1 & RIP v2 updates to the broadcast address.

timers

Syntax	timers <i>update timeout flush</i> no timers
Context	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor config>service>vprn>ripng config>service>vprn>ripng>group config>service>vprn>ripng>group>neighbor
Description	This command sets the values for the update, timeout, and flush timers. • Update timer — Determines how often RIP updates are sent. • Timeout timer — If a router is not updated by the time the timer expires, the route is declared invalid, but maintained in the RIP database. • Flush timer — Determines how long a route is maintained in the RIP database, after it has been declared invalid. Once this timer expires it is flushed from the RIP database completely. The no form of the command resets all timers to their default values of 30, 180, and 120 seconds respectively.
Default	no timers
Parameters	<i>update</i> — The RIP update timer value in seconds. Values 1 to 600 Default 30 <i>timeout</i> — The RIP timeout timer value in seconds. Values 1 to 1200 Default 180 <i>flush</i> — The RIP flush timer value in seconds. Values 1 to 1200 Default 120

unicast-address

Syntax	[no] unicast-address <i>ipv6-address</i>
Context	config>service>vprn>ripng>group>neighbor
Description	This command configures the unicast IPv6 address, RIPng updates messages will be sent to if the RIPng send command is set to send unicast .

Multiple unicast-address entries can be configured, in which case unicast messages will be sent to each configured unicast IPv6 address.

The **no** form of the command deletes the specified IPv6 unicast address from the configuration.

group

Syntax	[no] group <i>group-name</i>
Context	config>service>vprn>rip config>service>vprn>ripng
Description	This command creates a context for configuring a RIP group of neighbors. RIP groups are a way of logically associating RIP neighbor interfaces to facilitate a common configuration for RIP interfaces. The no form of the command deletes the RIP neighbor interface group. Deleting the group will also remove the RIP configuration of all the neighbor interfaces currently assigned to this group.
Default	no group
Parameters	<i>group-name</i> — The RIP group name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

neighbor

Syntax	[no] neighbor <i>ip-int-name</i>
Context	config>service>vprn>rip>group config>service>vprn>ripng>group
Description	This command creates a context for configuring a RIP neighbor interface. By default, interfaces are not activated in any interior gateway protocol such as RIP unless explicitly configured. The no form of the command deletes the RIP interface configuration for this interface. The shutdown command in the config>router>rip>group <i>group-name</i>>neighbor <i>ip-int-name</i> context can be used to disable an interface without removing the configuration for the interface.
Default	no neighbor

Parameters *ip-int-name* — The IP interface name. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service vprn interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

If the IP interface name does not exist or does not have an IP address configured an error message will be returned.

3.8.2.37 SDP Commands

spoke-sdp

Syntax `[no] spoke-sdp sdp-id`

Context `config>service>vprn`

Description This command binds a service to an existing Service Distribution Point (SDP). A spoke SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a VPRN service. If the SDP-ID is not already configured, an error message is generated. If the SDP-ID does exist, a binding between that SDP-ID and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Special Cases **VPRN** — Several SDPs can be bound to a VPRN service. Each SDP must be destined to a different router. If two SDP-ID bindings terminate on the same router, an error occurs and the second SDP binding is rejected.

Parameters *sdp-id* — The SDP identifier. Allowed values are integers in the range of 1 and 17407 for existing SDPs.

vc-id — The virtual circuit identifier.

Values 1 to 4294967295

control-channel-status

Syntax	[no] control-channel-status
Context	config>service>vprn>spoke-sdp
Description	This command enables the configuration of static pseudowire status signaling on a spoke-SDP for which signaling for its SDP is set to OFF. A control-channel-status no shutdown is allowed only if all of the following are true: • SDP signaling is off. • The control-word is enabled (the control-word is disabled by default) • The service type is Epipe, Apipe, VPLS, Cpipe, or IES/VPRN • Mate SDP signaling is off (in vc-switched services) • The pw-path-id is configured for this spoke-SDP. The no form of this command removes control channel status signaling from a spoke-SDP. It can only be removed if control channel status is shut down.
Default	no control-channel-status

acknowledgment

Syntax	[no] acknowledgment
Context	config>service>vprn>spoke-sdp>control-channel-status
Description	This command enables the acknowledgment of control channel status messages. By default, no acknowledgment packets are sent.

refresh-timer

Syntax	refresh-timer <i>value</i> no refresh-timer				
Context	config>service>vprn>spoke-sdp>control-channel-status				
Description	This command configures the refresh timer for control channel status signaling packets. By default, no refresh packets are sent.				
Default	no refresh-timer				
Parameters	<i>value</i> — Specifies the refresh timer value. <table><tr><td>Values</td><td>10 to 65535 seconds</td></tr><tr><td>Default</td><td>0 (off)</td></tr></table>	Values	10 to 65535 seconds	Default	0 (off)
Values	10 to 65535 seconds				
Default	0 (off)				

request-timer

Syntax	request-timer <i>timer1</i> retry-timer <i>timer2</i> timeout-multiplier <i>multiplier</i> no request-timer
Context	config>service>vprn>spoke-sdp>control-channel-status
Description	This command configures the control channel status request mechanism. When it is configured, control channel status request procedures are used. These augment the procedures for control channel status messaging from RFC 6478. This command is mutually exclusive with a non-zero refresh-timer value.
Parameters	<i>timer1</i> — Specifies the interval at which pseudowire status messages, including a reliable delivery TLV, with the “request” bit set, are sent. Values 10 to 65535 seconds retry-timer <i>timer2</i> — Specifies the timeout interval if no response to a pseudowire status request is received. This parameter must be configured. A value of zero (0) disables retries. Values 0, 3 to 60 seconds timeout-multiplier <i>multiplier</i> — If a requesting node does not receive a valid response to a pseudowire status request within this multiplier times the retry timer, then it will assume the pseudowire is down. This parameter is optional. Values 3 to 20 seconds

control-word

Syntax	[no] control-word
Context	config>service>vprn>spoke-sdp
Description	The control word command provides the option to add a control word as part of the packet encapsulation for pseudowire types for which the control word is optional. These are Ethernet pseudowires (Epipe). ATM N:1 cell mode pseudowires (apipe vc-types atm-vcc and atm-vpc) and VT pseudowire (apipe vc-type atm-cell). The configuration for the two directions of the pseudowire must match because the control word negotiation procedures described in Section 6.2 of RFC 4447 are not supported. The C-bit in the pseudowire FEC sent in the label mapping message is set to 1 when the control word is enabled. Otherwise, it is set to 0.

The service will only come up if the same C-bit value is signaled in both directions. If a spoke-sdp is configured to use the control word but the node receives a label mapping message with a C-bit clear, the node releases the label with the an “Illegal C-bit” status code as per Section 6.1 of RFC 4447. As soon as the user also enabled the control the remote peer, the remote peer will withdraw its original label and will send a label mapping with the C-bit set to 1 and the VLL service will be up in both nodes. The control word must be enabled to allow MPLS-TP OAM to be used on a static spoke-sdp in a apipe, epipe and cpipe service.

pw-path-id

Syntax	[no] pw-path-id
Context	config>service>vprn>spoke-sdp
Description	This command enters the context to configure an MPLS-TP Pseudowire Path Identifier for a spoke-sdp. All elements of the PW path ID must be configured in order to enable a spoke-sdp with a PW path ID.

For an IES or VPRN spoke-sdp, the pw-path-id is only valid for ethernet spoke-sdps.

The **pw-path-id** is only configurable if all of the following is true:

- SDP signaling is off
- control-word is enabled (control-word is disabled by default)
- the service type is epipe, vpls, cpipe, apipe, or IES/VPRN interface
- mate SDP signaling is off for vc-switched services

The **no** form of the command deletes the PW path ID.

Default	no pw-path-id

agi

Syntax	agi agi no agi
Context	config>service>vprn>spoke-sdp>pw-path-id
Description	This command configures the attachment group identifier for an MPLS-TP PW.
Parameters	agi — Specifies the attachment group identifier. Values 0 to 4294967295

saii-type2

Syntax	saii-type2 <i>global-id:node-id:ac-id</i> no saii-type2
Context	config>service>vprn>spoke-sdp>pw-path-id
Description	This command configures the source individual attachment identifier (SAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured (for example, it is at an S-PE), then the values must match those of the taii-type2 of the mate spoke-sdp.
Parameters	<i>global-id</i> — Specifies the global ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. Values 0 to 4294967295 <i>node-id</i> — Specifies the node ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. Values a.b.c.d or 0 to 4294967295 <i>ac-id</i> — Specifies the attachment circuit ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value. Values 1 to 4294967295

taii-type2

Syntax	taii-type2 <i>global-id:node-id:ac-id</i> no taii-type2
Context	config>service>vprn>spoke-sdp>pw-path-id
Description	This command configures the target individual attachment identifier (TAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured (for example, it is at an S-PE), then the values must match those of the saii-type2 of the mate spoke-sdp.
Parameters	<i>global-id</i> — Specifies the global ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP. Values 0 to 4294967295 <i>node-id</i> — Specifies the node ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP. Values a.b.c.d or 0 to 4294967295

ac-id — Specifies the attachment circuit ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value.

Values 1 to 4294967295

spoke-sdp

Syntax **spoke-sdp** *sdp-id* [:*vc-id*] *vc-type* {**ether** | **ipipe**} [**create**]
no spoke-sdp *sdp-id* [:*vc-id*] *vc-type* {**ether** | **ipipe**} [**create**]

Context config>service>vprn>if

Description This command binds a service to an existing Service Distribution Point (SDP).

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

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a service. If the SDP-ID is not already configured, an error message is generated. If the SDP ID does exist, a binding between that SDP ID and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

Class-based forwarding is not supported on a spoke SDP used for termination on an IES or VPRN services. All packets are forwarded over the default LSP.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once 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 7750 SR OS. If two SDP-ID bindings terminate on the same 7750 SR, 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.

Values 1 to 4294967295

vc-type — The encapsulation and pseudowire type for the spoke SDP.

Values ether—Ethernet pseudowire.

ipipe—IP pseudowire.

Default ether

egress

Syntax	egress
Context	config>service>vprn>if>spoke-sdp config>service>vprn>red-if>spoke-sdp
Description	This command configures an SDP context.

entropy-label

Syntax	[no] entropy-label
Context	config>service>vprn config>service>vprn>if>spoke-sdp
Description	This command enables or disables the use of entropy labels for spoke SDPs on a VPRN. If entropy-label is configured, the entropy label and ELI are inserted in packets for which at least one LSP in the stack for the far-end of the tunnel used by the service has advertised entropy-label-capability. If the tunnel is RSVP type, entropy-label can also be controlled under the config>router>mpls or config>router>mpls>lsp contexts. The entropy label and the hash label features are mutually exclusive. The entropy label cannot be configured on a spoke SDP or service where the hash label feature has already been configured.
Default	no entropy-label

hash-label

Syntax	[no] hash-label
Context	config>service>vprn config>service>vprn>spoke-sdp config>service>vprn>if>spoke-sdp
Description	This command enables the use of the hash label on a VLL, VPLS, or VPRN service bound to any MPLS-type encapsulated SDP as well as to a VPRN service using auto-bind-tunnel with the resolution-filter configured as any MPLS tunnel type. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the gre option.

When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to 1 to indicate that.

In order to allow for applications whereby 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 will always be in the range [524,288 - 1,048,575] and will not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label will not match a value in the reserved label range.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. For VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL pseudowire packets.

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

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

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

Default no hash-label

ingress

Syntax ingress

Context config>service>vprn>if>spoke-sdp
config>service>vprn>red-if>spoke-sdp

Description This command configures the SDP context.

qos

Syntax qos *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*
no qos

Context config>service>vprn>if>spoke-sdp>ingress

Description This command is used to redirect pseudowire packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.

The ingress pseudowire rate-limiting feature uses a policer in queue-group provisioning model. This model allows the mapping of one or more pseudowires to the same instance of policers, which are defined in a queue-group template.

Operationally, the provisioning model in the case of the ingress pseudowire shaping feature consists of the following steps:

- Step 1.** Create an ingress queue-group template and configure policers for each FC that needs to be redirected and optionally, for each traffic type (unicast, broadcast, unknown, or multicast).
- Step 2.** Apply the queue-group template to the network ingress forwarding plane where there exists a network IP interface to which the pseudowire packets can be received. This creates one instance of the template on the ingress of the FP. One or more instances of the same template can be created.
- Step 3.** Configure FC-to-policer mappings together with the policer redirect to a queue-group in the ingress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different pseudowires to different queue-group templates.
- Step 4.** Apply this network QoS policy to the ingress context of a spoke-SDP inside a service, or to the ingress context of a pseudowire template, and specify the redirect queue-group name.
- Step 5.** One or more spoke-SDPs can have their FCs redirected to use policers in the same policer queue-group instance.

The following are the constraints and rules of this provisioning model when used in the ingress pseudowire rate-limiting feature:

- Step 1.** When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name does not exist, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
- Step 2.** When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists but the policer-id is not defined in the queue-group template, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
- Step 3.** When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists and the policer-id is defined in the queue-group template, it is not required to check that an instance of that queue-group exists in all ingress FPs which have network IP interfaces. The handling of this is dealt with in the data path as follows:
 - a. When a pseudowire packet for that FC is received and an instance of the referenced queue-group name exists on that FP, the packet is processed by the policer and will then feed the per-FP ingress shared queues referred to as *policer-output-queues*.

- b. When a pseudowire packet for that FC is received and an instance of the referenced queue-group name does not exist on that FP, the pseudowire packets will be fed directly into the corresponding ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.

Step 4. If a network QoS policy is applied to the ingress context of a pseudowire, any pseudowire FC which is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.

- Step 5.** If no network QoS policy is applied to the ingress context of the pseudowire, all packets of the pseudowire will feed:
- a. the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the MDA/FP. This is the default behavior.
 - b. a queue-group policer followed by the per-FP ingress shared queues referred to as *policer-output-queues* if the ingress context of the network IP interface from which the packet is received is redirected to a queue-group (csc-policing). The only exceptions to this behavior are for packets received from a IES/VPRN spoke interface and from an R-VPLS spoke-SDP, which is forwarded to the R-VPLS IP interface. In these two cases, the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the MDA/FP is used.

When a pseudowire is redirected to use a policer queue-group, the classification of the packet for the purpose of FC and profile determination is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the pseudowire. This is true regardless of whether an instance of the named policer queue-group exists on the ingress FP on which the pseudowire packet is received. The user can apply a QoS filter matching the dot1p in the VLAN tag corresponding to the Ethernet port encapsulation, the EXP in the outer label when the tunnel is an LSP, the DSCP in the IP header if the tunnel encapsulation is GRE, and the DSCP in the payload IP header if the user enabled the **ler-use-dscp** option and the pseudowire terminates in IES or VPRN service (spoke-interface).

When the policer queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the packet classification is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface on which the pseudowire packet is received.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

Parameters *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 to 65535

fp-redirect-group *queue-group-name* — Specifies the name of the queue group template up to 32 characters in length.

ingress-instance *instance-id* — Specifies the identification of a specific instance of the queue-group.

Values 1 to 16384

vc-label

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

Context config>service>vprn>if>spoke-sdp>egress
config>service>vprn>red-if>spoke-sdp>egress

Description This command configures the egress VC label.

Parameters *vc-label* — A VC egress value that indicates a specific connection.

Values 16 to 1048575

vc-label

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

Context config>service>vprn>if>spoke-sdp>ingress
config>service>vprn>red-if>spoke-sdp>ingress

Description This command configures the ingress VC label.

Parameters *vc-label* — A VC ingress value that indicates a specific connection.

Values 2048 to 18431

egress

Syntax **egress**

Context config>service>vprn>network-interface

Description This command enters the context to configure egress network filter policies for the interface.

filter

Syntax **filter ip** *ip-filter-id*
filter ipv6 *ipv6-filter-id*
no filter [*ip ip-filter-id*] [*ipv6 ipv6-filter-id*]

Context	<pre> config>service>vprn>nw-if>egress config>service>vprn>if>spoke-sdp>egress config>service>vprn>if>spoke-sdp>ingress config>service>vprn>red-if>spoke-sdp>ingress config>service>vprn>red-if>spoke-sdp>egress config>service>vprn>nw-if>egress </pre>
Description	This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface. An IP filter policy can be associated with spoke SDPs. Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria. The filter command is used to associate a filter policy with a specified ip-filter-id with an ingress or egress SAP. The ip-filter-id must already be defined before the filter command is executed. If the filter policy does not exist, the operation fails and an error message returned. In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets. The no form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use scope command within the filter definition to change the scope to local or global. The default scope of a filter is local.
Parameters	ip <i>ip-filter-id</i> — Specifies IP filter policy. The filter ID must already exist within the created IP filters. Values 1 to 65535

qos

Syntax	<pre> qos <i>network-policy-id</i> port-redirect-group <i>queue-group-name</i> [instance <i>instance-id</i>] no qos [<i>network-policy-id</i>] </pre>
Context	<pre> config>service>pw-template>egress config>service>vprn>if>spoke-sdp>egress config>service>ies>if>spoke-sdp>egress </pre>
Description	This command is used to redirect pseudowire packets to an egress port queue-group for the purpose of shaping. The egress pseudowire shaping provisioning model allows the mapping of one or more pseudowires to the same instance of queues, or policers and queues, which are defined in the queue-group template. Operationally, the provisioning model consists of the following steps: Step 1. Create an egress queue-group template and configure queues only or policers and queues for each FC that needs to be redirected.

- Step 2.** Apply the queue-group template to the network egress context of all ports where there exists a network IP interface on which the pseudowire packets can be forwarded. This creates one instance of the template on the egress of the port. One or more instances of the same template can be created.
- Step 3.** Configure FC-to-policer or FC-to-queue mappings together with the redirect to a queue-group in the egress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different pseudowires to different queue-group templates.
- Step 4.** Apply this network QoS policy to the egress context of a spoke-SDP inside a service or to the egress context of a pseudowire template and specify the redirect queue-group name.

One or more spoke-SDPs can have their FCs redirected to use queues only or queues and policers in the same queue-group instance.

The following are the constraints and rules of this provisioning model:

- Step 1.** When a pseudowire FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name does not exist, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface on which the pseudowire packet is forwarded. This queue can be a queue-group queue, or the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port. This is the existing implementation and default behavior for a pseudowire packet.
- Step 2.** When a pseudowire FC is redirected to use a queue or a policer, and a queue in a queue-group and the queue-group name exists, but the policer-id and/or the queue-id is not defined in the queue-group template, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the pseudowire packet is forwarded on.
- Step 3.** When a pseudowire FC is redirected to use a queue, or a policer and a queue in a queue-group, and the queue-group name exists and the policer-id or policer-id plus queue-id exist, it is not required to check that an instance of that queue-group exists in all egress network ports which have network IP interfaces. The handling of this is dealt with in the data path as follows:
 - a. When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name exists on that egress port, the packet is processed by the queue-group policer and will then be fed to the queue-group queue.
 - b. When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name does not exist on that egress port, the pseudowire packet will be fed directly to the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

Step 4. If a network QoS policy is applied to the egress context of a pseudowire, any pseudowire FC, which is not explicitly redirected in the network QoS policy, will have the corresponding packets feed directly the corresponding the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

When the queue-group name the pseudowire is redirected to exists and the redirection succeeds, the marking of the packet DEI/dot1p/DSCP and the tunnel DEI/dot1p/DSCP/EXP is performed; according to the relevant mappings of the (FC, profile) in the egress context of the network QoS policy applied to the pseudowire. This is true regardless, whether an instance of the queue-group exists or not on the egress port to which the pseudowire packet is forwarded. If the packet profile value changed due to egress child policer CIR profiling, the new profile value is used to mark the packet DEI/dot1p and the tunnel DEI/dot1p/EXP, and the DSCP/prec will be remarked if **enable-dscp-prec-marking** is enabled under the policer.

When the queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the marking of the packet DEI/dot1p/DSCP and the tunnel DEI/dot1p/DSCP/EXP fields is performed according to the relevant commands in the egress context of the network QoS policy applied to the network IP interface to which the pseudowire packet is forwarded.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

Parameters *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1 to 65535

port-redirect-group *queue-group-name* — This optional parameter specifies that the *queue-group-name* will be used for all egress forwarding class redirections within the network QoS policy ID. The specified *queue-group-name* must exist as a port egress queue group on the port associated with the IP interface.

egress-instance *instance-id* — Specifies the identification of a specific instance of the queue-group.

Values 1 to 16384

3.8.2.38 RADIUS Proxy Commands

radius-proxy

Syntax	radius-proxy
Context	config>service>vprn
Description	This command enters the context to configure RADIUS proxy commands.

server

Syntax	server <i>server-name</i> [create] [purpose {[accounting] [authentication]}] no server <i>server-name</i>
Context	config>service>vprn>radius-proxy
Description	This command configures the name of this RADIUS proxy server.
Parameters	purpose accounting — Specifies that this RADIUS proxy server will be used for accounting purposes. purpose authentication — Specifies that this RADIUS proxy server will be used for authentication purposes.

cache

Syntax	cache
Context	config>service>vprn>radius-proxy>server
Description	This command enters the context to configure caching parameters. The no form of the command disables caching.

key

Syntax	key <i>packet-type</i> { accept request } attribute-type <i>attribute-type</i> [vendor-id <i>vendor-id</i>] no key
Context	config>service>vprn>radius-proxy>server>cache
Description	This command configures cache key parameters. The no form of the command removes the parameters from the configuration.
Default	no key
Parameters	<i>packet-type</i> — Specifies the packet type of the RADIUS messages to use to generate the key for the cache of this RADIUS proxy server. In order to generate the key associated with a RADIUS Access-Accept message, the system uses the attribute of the type specified by the value of <code>tmnxRadProxSrvCacheKeyAttrType</code> , within the associated RADIUS message of the type specified by the value of <code>tmnxRadProxSrvCacheKeyPktType</code> . Values accept, request attribute-type <i>attribute-type</i> — Specifies the RADIUS attribute type to cache for this RADIUS Proxy server.

In order to generate the key associated with a RADIUS Access-Accept message, the system uses the attribute of the type specified by the value of `tmnxRadProxSrvCacheKeyAttrType`, within the associated RADIUS message of the type specified by the value of `tmnxRadProxSrvCacheKeyPktType`.

Values 1 to 255

vendor-id *vendor-id* — Specifies the RADIUS Vendor-Id.

If the value of `tmnxRadProxSrvCacheKeyVendorId` is equal to zero, the attribute type specified by `tmnxRadProxSrvCacheKeyAttrType` must be used if it appears outside of a Vendor-Specific attribute.

If the value of `tmnxRadProxSrvCacheKeyVendorId` is not equal to zero, the attribute type specified by `tmnxRadProxSrvCacheKeyAttrType` must be used if it appears as a sub-attribute within a Vendor-Specific attribute with Vendor-Id equal to the value of `tmnxRadProxSrvCacheKeyVendorId`.

Values 1 to 16777215, alu

timeout

Syntax	timeout [<i>hrs hours</i>] [<i>min minutes</i>] [<i>sec seconds</i>] no timeout
Context	config>service>vprn>radius-proxy>server>cache
Description	This command configures the timeout, in seconds, after which an entry in the cache will expire. The no form of the command reverts to the default.
Default	timeout min 5
Parameters	<i>timeout</i> — Configures the timeout. Values hours: 1 to 1 minutes: 1 to 59 seconds: 1 to 59

track-accounting

Syntax	track-accounting [<i>start</i>] [<i>stop</i>] [<i>interim-update</i>] no track-accounting
Context	config>service>vprn>radius-proxy>server>cache
Description	This command specifies which RADIUS accounting packets have impact on the cache of this RADIUS proxy server. Use it to configure what RADIUS accounting packets have impact on the cache.

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

Default no track-accounting

default-accounting-server-policy

Syntax **default-accounting-server-policy** *policy-name*
no default-accounting-server-policy

Context config>service>vprn>radius-proxy>server

Description This command configures the name of the default RADIUS server policy associated with this RADIUS proxy server for accounting purposes.

This default policy is used if no policy can be derived from the user name.

The **no** form of the command removes the policy from the configuration.

Parameters *policy-name* — Specifies the default accounting RADIUS server policy up to 32 characters in length.

default-authentication-server-policy

Syntax **default-authentication-server-policy** *policy-name*
no default-authentication-server-policy

Context config>service>vprn>radius-proxy>server

Description This command configures the name of the default RADIUS server policy associated with this RADIUS proxy server for authentication purposes.

This default policy is used if no policy can be derived from the user name.

The **no** form of the command removes the policy from the configuration.

Parameters *policy-name* — Specifies the default authentication RADIUS server policy up to 32 characters in length.

interface

Syntax [**no**] **interface** *ip-int-name*

Context config>service>vprn>radius-proxy>server

Description This command associates an interface to the proxy server.

The no form of the command removes the interface name from the proxy server configuration.

Parameters *ip-int-name* — Specifies the name of the IP interface. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed between double quotes.

load-balance-key

Syntax **load-balance-key** *vendor vendor-id* [*vendor-id...(up to 5 max)*] **attribute-type** *attribute-type* [*attribute-type...(up to 5 max)*]
load-balance-key *source-ip-udp*
no load-balance-key

Context config>service>vprn>radius-proxy>server

Description This command configures how to construct the key for load-balancing RADIUS messages between RADIUS servers.

Default no load-balance-key

Parameters **vendor** *vendor-id* — Specifies the RADIUS Vendor-Id.
Values 0 to 16777215
attribute-type *attribute-type* — Specifies a RADIUS attribute that must be used to construct the key for load-balancing RADIUS messages between RADIUS servers.
Values 1 to 255

python-policy

Syntax **python-policy** *name*
no python-policy

Context config>service>vprn>radius-proxy>server
config>service>vprn>radius-server>server

Description This common specifies a python policy. Python policies are configured in the **config>python> python-policy** *name* context.

Parameters *name* — Specifies the name of an existing python script up to 32 characters in length.

secret

Syntax **secret** *secret* [*hash* | *hash2*]
no secret

Context	config>service>vprn>radius-proxy>server
Description	This command configures the secret key associated with the RADIUS server. The no form of the command removes the key from the configuration.
Parameters	secret — The secret key (password) to access the RADIUS server. This secret key must match the password on the RADIUS server. Values <i>secret-key</i>: Up to 20 characters in length. <i>hash-key</i>: Up to 33 characters in length. <i>hash2-key</i>: Up to 55 characters in length. hash — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified. hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the hash2 encrypted variable cannot be copied and pasted. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.

send-accounting-response

Syntax	[no] send-accounting-response
Context	config>service>vprn>radius-proxy>server
Description	This command specifies if this RADIUS proxy server itself responds with an Accounting-Response message to each received Accounting-Request instead of proxying them to a configured RADIUS server. The no form of the command disables the accounting response messages.
Default	no send-accounting-response

username

Syntax	username <i>user-name</i> prefix-string [128 chars max] [accounting-server-policy <i>policy-name</i>] [authentication-server-policy <i>policy-name</i>] no username [1 to 32]
Context	config>service>vprn>radius-proxy>server
Description	This command configures username-to-RADIUS-server-policy associations. The no form of the command removes the associations from the configuration.

Parameters	<i>username</i> — Specifies the user name.
Values	1 to 32
prefix-string	— Specifies the prefix-string for the association.
Values	128 characters, maximum
accounting-server-policy	<i>policy-name</i> — Specifies the default accounting RADIUS server policy up to 32 characters in length.
authentication-server-policy	<i>policy-name</i> — Specifies the default authentication RADIUS server policy up to 32 characters in length.

radius-server

Syntax	radius-server
Context	config>service>vprn
Description	This command enters the context to configure RADIUS server parameters.

server

Syntax	server <i>server-name</i> [address <i>ip-address</i>] [secret <i>key</i>] [hash hash2] [port <i>port</i>] [create] no server <i>server-name</i>
Context	config>service>vprn>radius-server
Description	This command configures RADIUS server parameters. The no form of the command removes the parameters from the configuration.
Parameters	<i>server-name</i> — Specifies the name of this RADIUS server. address <i>ip-address</i> — Specifies the IP address of the RADIUS server. secret <i>key</i> — Specifies the secret key (password) to access the RADIUS server. This secret key must match the password on the RADIUS server. Values <i>secret-key</i> : Up to 20 characters in length. <i>hash-key</i> : Up to 33 characters in length. <i>hash2-key</i> : Up to 55 characters in length. hash — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

port *port* — Specifies the UDP port number on which to contact the RADIUS server.

accept-coa

Syntax	[no] accept-coa
Context	config>service>vprn>radius-server
Description	This command specifies if this RADIUS server is allowed to process Change of Authorization messages.

coa-script-policy

Syntax	coa-script-policy <i>script-policy-name</i> no coa-script-policy
Context	config>service>vprn>radius-server
Description	This command specifies the RADIUS script policy used to change the RADIUS attributes of the Change-of-Authorization messages. The no form of the command removes the script policy from the configuration.
Parameters	<i>script-policy-name</i> — Specifies a Python script policy to modify Change-of-Authorization messages.

pending-requests-limit

Syntax	pending-requests-limit <i>limit</i> no pending-requests-limit
Context	config>router config>service>vprn>radius-server>server
Description	This command specifies the limit of the number of pending RADIUS authentication requests.
Default	pending-requests-limit 4096
Parameters	<i>limit</i> — Configure the limit of the number of pending RADIUS requests. Values 1 to 4096

wpp

Syntax	[no] wpp
Context	config>router config>service>vprn
Description	This command enters the configuration context of web portal protocol (WPP) under router or vprn. The no form of this command removes configuration under wpp.
Default	no wpp

portals

Syntax	portals
Context	config>router>wpp config>service>vprn>wpp
Description	This command enters the configuration context of web portal server.

portal

Syntax	portal <i>name</i> address <i>ip-address</i> [create] [no] portal <i>name</i>
Context	config>router>wpp>portals config>service>vprn>wpp>portals
Description	This command either creates a new web portal server or enters an existing web portal server.
Default	no portal
Parameters	<i>name</i> — Specifies the name of the web portal server. <i>ip-address</i> — Specifies IPv4 address of the web portal server.

wpp

Syntax	[no] wpp
Context	config>service>ies>sub-if>grp-if> config>service>vprn>sub-if>grp-if>
Description	This command enters the configuration context of web portal protocol (WPP) under group-interface.

The **no** form of this command removes configuration under WPP.

Default no wpp

initial-app-profile

Syntax	initial-app-profile <i>profile-name</i> no initial-app-profile
Context	config>service>ies>sub-if>grp-if>wpp config>service>vprn>sub-if>grp-if>wpp
Description	This command specifies the initial app-profile for the hosts created on the group-interface. This initial app-profile is replaced after hosts pass the web portal authentication.
Default	no initial-app-profile
Parameters	<i>profile-name</i> — Specifies the name of app-profile.

initial-sla-profile

Syntax	initial-sla-profile <i>profile-name</i> no initial-sla-profile
Context	config>router>wpp config>service>vprn>wpp
Description	This command specifies the initial sla-profile for the hosts created on the group-interface. This initial sla-profile is replaced after hosts pass the web portal authentication.
Default	no initial-sla-profile
Parameters	<i>profile-name</i> — Specifies the name of sla-profile.

initial-sub-profile

Syntax	initial-sub-profile <i>profile-name</i> no initial-sub-profile
Context	config>service>ies>sub-if>grp-if>wpp config>service>vprn>sub-if>grp-if>wpp
Description	This command specifies the initial sub-profile for the hosts created on the group-interface. This initial sub-profile will be replaced after hosts pass web portal authentication.
Default	no initial-sub-profile

Parameters *profile-name* — Specifies the name of sub-profile.

portal

Syntax **portal router** *router-instance name wpp-portal-name*
no portal

Context config>service>ies>sub-if>grp-if>wpp
 config>service>vprn>sub-if>grp-if>wpp

Description This command specifies the web portal server that system talks to for the hosts on the group-interface.

Default no portal

Parameters *router-instance* — Specifies the routing-instance that web portal server is defined.
 profile-name — Specifies the name of the web portal server.

restore-disconnected

Syntax [no] **restore-disconnected**

Context config>service>ies>sub-if>grp-if>wpp
 config>service>vprn>sub-if>grp-if>wpp

Description This command enable the behavior that system will restore the **initial-sla-profile**, **initial-sub-profile**, or **initial-app-profile** when hosts disconnects instead of removing them.

 The **no** form of the command specifies that the initial profiles will not be restored after a DHCP host has disconnected.

Default restore-disconnected

3.8.2.39 AARP Interface Commands

aarp-interface

Syntax **aarp-interface** *aarp-interface-name* [create]
no aarp-interface *aarp-interface-name*

Context config>service>vprn

Description	This command creates an AARP interface for connecting a service to a peer node AARP service. This instance is paired with the same AARP interface in a peer node as part of a configuration to provide flow and packet asymmetry removal for traffic for a multi-homed SAP or spoke SDP. The no form of the command deletes the interface.
Default	no aarp-interface
Parameters	<i>aarp-interface-name</i> — A string of up to 32 characters identifying the interface. create — Keyword used to create the AARP interface.

ip-mtu

Syntax	ip-mtu <i>octets</i> no ip-mtu
Context	config>service>vprn>aarp-interface
Description	This command configures the IP maximum transmit unit (packet) for this interface. The no form of the command returns the default value. By default (for Ethernet network interface) if no ip-mtu is configured it is (1568 - 14) = 1554.
Default	no ip-mtu
Parameters	<i>octets</i> — Specifies the maximum number of octets that can be transmitted. Values 512 to 9000

spoke-sdp

Syntax	spoke-sdp <i>sdp-id:vc-id</i> [create] no spoke-sdp <i>sdp-id:vc-id</i>
Context	config>service>vprn>aarp-interface
Description	This command binds a service to an existing SDP. A spoke SDP is treated like the equivalent of a traditional bridge port where flooded traffic received on the spoke SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received. The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down. SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Default	no spoke-sdp
Parameters	<i>sdp-id</i> — Specifies the SDP identifier. Values 1 to 17407 <i>vc-id</i> — The virtual circuit identifier. The VC-ID is not used with L2TPv3 SDPs, however it must be configured. Values 1 to 4294967295 create — Keyword used to create the spoke SDP.

aarp

Syntax	aarp <i>aarp-id</i> type {subscriber-side-shunt network-side-shunt} no aarp
Context	config>service>vprn>aarp-interface>spoke-sdp
Description	This command associates an AARP instance to an AARP interface spoke SDP. This instance is paired with the same <i>aarp-id</i> in a peer node as part of a configuration to provide flow and packet asymmetry removal for traffic for a multi-homed SAP or spoke SDP. The type parameter specifies the role of this service point in the AARP instance. The no form of the command removes the association.
Default	no aarp
Parameters	<i>aarp-id</i> — An integer that identifies an AARP instance. Values 1 to 65535 subscriber-side-shunt — Specifies that the AARP type is an inter-chassis shunt service for subscriber-side traffic. network-side-shunt — Specifies that the AARP type is an inter-chassis shunt service for network-side traffic.

egress

Syntax	egress
Context	config>service>vprn>aarp-interface>spoke-sdp
Description	This command enters the egress context for a spoke SDP.

filter

Syntax	filter ip <i>ip-filter-id</i> no filter
Context	config>service>vprn>aarp-interface>spoke-sdp>egress config>service>vprn>aarp-interface>spoke-sdp>ingress
Description	This command associates an IP filter policy with an ingress or egress IP interface. Filter policies control the forwarding and dropping of packets based on IP matching criteria. The <i>filter-id</i> must already be defined before the filter command is executed. If the filter policy does not exist, the operation fails and an error message returned. IP filters apply only to RFC 2427-routed IP packets. Frames that do not contain IP packets will not be subject to the filter and will always be passed, even if the filter's default action is to drop. The no form of this command removes any configured filter ID association with the IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local.
Parameters	<i>ip-filter-id</i> — Specifies the filter policy. The filter ID must already exist within the created IP filters. Values 1 to 65535 or a string up to 64 characters

vc-label

Syntax	vc-label <i>vc-label</i> no vc-label [<i>vc-label</i>]
Context	config>service>vprn>aarp-interface>spoke-sdp>egress config>service>vprn>aarp-interface>spoke-sdp>ingress
Description	This command configures the egress and ingress VC label. The no version of this command removes the VC label.
Parameters	<i>vc-label</i> — A VC egress value that indicates a specific connection. Values egress: 16 to 1048575 ingress: 32 to 18431

ingress

Syntax	ingress
Context	config>service>vprn>aarp-interface>spoke-sdp

Description This command enters the ingress context for a spoke SDP.

3.9 VPRN Show, Clear, and Debug Command Reference

3.9.1 Command Hierarchies

- [Show Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)

3.9.1.1 Show Commands

```

show
  — service
    — egress-label start-label [end-label]
    — id service-id
      — all
      — arp [ip-address | mac ieee-address | sap port-id:encap | interface ip-int-name
        | sdp sdp-id:vc-id | summary]
      — arp-host [wholesaler service-id] [sap sap-id | interface interface-name | ip-
        address ip-address[/mask] | mac ieee-address | {[port port-id] [no-inter-
        dest-id | inter-dest-id inter-dest-id]]} [detail]
      — arp-host statistics [sap sap-id | interface interface-name]
      — arp-host summary [interface interface-name | saps]
      — authentication
        — statistics [policy name] [sap sap-id]
      — base [msap]
      — dhcp
        — lease-state [wholesaler service-id] [sap sap-id | sdp sdp-id:vc-id |
          interface interface-name | ip-address ip-address[/mask] | chaddr
          ieee-address | mac ieee-address | {[port port-id] [no-inter-dest-id
          inter-dest-id]]} [session {none | ipoe}] [detail]
        — statistics [sap sap-id | sdp sdp-id:vc-id | interface interface-name]
        — summary [interface interface-name | saps]
      — gsmpp
        — neighbors [group name [ip-address]]
        — sessions {group name | neighbor ip-address port port-number
          [association | statistics]}
      — interface {[ip-address | ip-int-name] [interface-type] [detail] [family]} |
        summary
      — ptp
      — retailers
      — sap sap-id static-isids [range-id range-id]
      — sap sap-id dist-cpu-protection [detail]
      — sap sap-id detail
  
```

```

— sap sap-id encap-group group-name [member encap-id] [encap-detail |
  encap-stats]
— sap sap-id encap-group
— sap sap-id [atm | base | dhcp | mc-ring | mcac | mrp | qos | sap-stats | stats
  | stp | sub-mgmt | ipsec-gw]
— sap sap-id [circuit-id circuit-id] [mac ieee-address] [remote-id remote-id]
  host-lockout-policy [summary]
— sap queue-depth [queue queue-id] [ingress | egress]
— sap
— sap sap-id static-isids mfib
— sap sap-id queue-group-redirect [ingress | egress]
— sdp [sdp-id [:vc-id]] [detail]
— sdp far-end {ip-address | ipv6-address} [detail]
— sdp sdp-id [:vc-id] l2tpv3
— sdp sdp-id [:vc-id] static-isids [range-id range-id]
— sdp sdp-id [:vc-id] static-isids mfib
— sdp sdp-id [:vc-id] [detail] vccv-bfd [session]
— sdp sdp-id [:vc-id] mrp
— subscriber-hosts [sap sap-id] [ip ip-prefix/prefix-length] [mac ieee-address]
  [sub-profile sub-profile-name] [sla-profile sla-profile-name] [app-profile
  app-profile-name] [wholesaler service-id] [address-origin address-origin]
  [detail]
— wholesalers
— ingress-label start-label [end-label]
— sap-using [msap] [dyn-script] [description]
— sap-using [sap sap-id] [vlan-translation | anti-spoof] [description]
— sap-using {ingress | egress} atm-td-profile td-profile-id
— sap-using {ingress | egress} filter any-filter-id
— sap-using {ingress | egress} qos-policy qos-policy-id [msap]
— sap-using etree
— sdp sdp-id pw-port [pw-port-id]
— sdp sdp-id pw-port
— sdp sdp-id pw-port pw-port-id [statistics]
— sdp [consistent | inconsistent | na] egressifs
— sdp sdp-id keep-alive-history
— sdp far-end {ip-address | ipv6-address} keep-alive-history
— sdp [sdp-id] [detail]
— sdp far-end {ip-address | ipv6-address} [detail]
— service-using [epipe] [ies] [vppls] [vprn] [mirror] [apipe] [fpipe] [ipipe] [cpipe] [etree]
  [vsd] [b-vpls] [i-vpls] [m-vpls] [sdp sdp-id] [customer customer-id] [origin origin-creation]

show
  — service
    — id service-id
      — log
        — filter-id [filter-id]
        — log-id [log-id] [severity severity-level] [application application]
          [sequence from-seq [to-seq]] [count count] [subject subject]
          [regex] [ascending | descending] [message format [msg-regex]]
        — snmp-trap-group [log-id]
        — syslog [syslog-id]

```



```

show
  — router [vprn-service-id]
    — aggregate [family] [active] [detail]
    — arp [ip-int-name | ip-address [/mask] | mac ieee-mac-address | summary] [arp-type]
    — bgp
      — damping [ip-prefix [/ip-prefix-length]] [damp-type] [detail] [ipv4]
      — damping [ip-prefix [/ip-prefix-length]] [damp-type] [detail] {ipv6 | label-ipv4 |
        label-ipv6 | mcast-ipv4 | mcast-ipv6 | mvpn-ipv4 | vpn-ipv4 | vpn-ipv6}
      — group [name] [detail]
      — neighbor [ip-address [detail]]
      — neighbor [as-number [detail]]
      — paths
      — routes [ip-prefix/mask | ip-address]
      — summary [all]
      — summary [family family] [group name]
    — dhcp
      — statistics [interface ip-int-name | ip-address]
      — summary
    — ecmp
    — ldp
      — bindings
    — mvpn
    — rip
      — database [ip-address[/mask] [longer]] [peer ip-address] [detail [qos]]
      — neighbor [ip-int-name | ip-address] [detail] [advertised-routes]
      — peer [interface-name]
      — statistics [ip-int-name | ip-address]
    — route-table [family] [ip-prefix[/prefix-length] [longer | exact | protocol protocol-name]
      [all] [next-hop-type type] [qos] [alternative] [accounting-class]
    — route-table [family] [summary]
    — route-table tunnel-endpoints [ip-prefix[/prefix-length] [longer | exact] [detail]
    — route-table [family] [ip-prefix[/prefix-length] [longer | exact | protocol protocol-name]
      extensive [all]
    — service-prefix
    — static-arp [ip-address | ip-int-name | mac ieee-mac-address]
    — static-route [family] [ip-prefix/prefix-length | [preference preference] | [next-hop ip-
      address | tag tag] [detail]
    — tunnel-table summary [ipv4 | ipv6]
    — tunnel-table [protocol protocol] {ipv4 | ipv6}
    — tunnel-table [ip-prefix[/mask]] [alternative] [ipv4 | ipv6] detail
    — tunnel-table [ip-prefix[/mask]] [alternative]
    — tunnel-table mpls-tp
    — tunnel-table [ip-prefix[/mask]] protocol protocol [detail]
    — tunnel-table [ip-prefix[/mask]] sdp sdp-id
    — wpp
    — wpp [portal wpp-portal-name] [host ip-address] hosts
    — wpp portal wpp-portal-name
    — wpp statistics
  
```

3.9.1.2 Clear Commands

```

clear
  — router
    — bgp
      — damping [ip-prefix/ip-prefix-length [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 |
        hard]
      — neighbor {ip-address | as as-number | external | all} statistic
      — neighbor ip-address end-of-rib
      — protocol
    — dhcp
      — statistics [ip-int-name | ip-address]
    — forwarding-table [slot-number]
    — interface [ip-int-name | ip-address] [urpf-stats] [statistics] [hold-time]
    — interface [ip-int-name | ip-address] [policy-accounting] [class] [index]
    — interface {ip-int-name | ip-address} mac [ieee-address]
    — rip
      — database
      — statistics [neighbor ip-int-name | ip-address]
      — statistics interface [ip-int-name | ip-address]

clear
  — service
    — id service-id
      — arp-host {all | mac ieee-address | sap sap-id | ip-address ip-address[/mask]}
      — arp-host {port port-id | {inter-dest-id intermediate-destination-id | no-inter-
        dest-id} [port port-id]}
      — arp-host statistics [sap sap-id | interface interface-name]
      — dhcp
        — lease-state all [no-dhcp-release]
        — lease-state [port port-id] inter-dest-id intermediate-destination-id [no-
          dhcp-release]
        — lease-state [port port-id] no-inter-dest-id [no-dhcp-release]
        — lease-state ip-address ip-address[/mask] [no-dhcp-release]
        — lease-state mac ieee-address [no-dhcp-release]
        — lease-state port port-id [no-dhcp-release]
        — lease-state sap sap-id [no-dhcp-release]
        — lease-state sdp sdp-id:vc-id [no-dhcp-release]
      — spoke-sdp sdp-id:vc-id ingress-vc-label
      — spoke-sdp sdp-id:vc-id l2tpv3
      — spoke-sdp sdp-id:vc-id vccv-bfd {session | statistics}
    — statistics
      — sap sap-id {all | counters | stp | l2pt | mrp}
      — sdp sdp-id keep-alive
      — id service-id
        — counters
        — spoke-sdp sdp-id:vc-id {all | counters | stp | l2pt | mrp}

```

3.9.1.3 Debug Commands

For protocol debugging commands, see the *7450 ESS, 7750 SR, 7950 XRS, and VSR Unicast Routing Protocols Guide* and *7450 ESS, 7750 SR, 7950 XRS, and VSR Multicast Routing Protocols Guide*.

```
debug
  — service
    — id service-id
      — [no] arp-host
      — [no] dhcp
        — detail-level {low | medium | high}
        — no detail-level
        — mode {dropped-only | ingr-and-dropped | egr-ingr-and-dropped}
        — no mode
        — [no] sap sap-id
        — [no] sdp sdp-id:vc-id
      — [no] event-type {config-change | svc-oper-status-change | sap-oper-
        status-change | sdpbind-oper-status-change}
      — [no] host-connectivity-verify
        — [no] ip ip-address
        — [no] mac ieee-address
        — [no] sap sap-id
      — [no] sap sap-id
        — event-type {arp | config-change | oper-status-change | neighbor-
          discovery}
      — [no] sdp sdp-id:vc-id
        — event-type {config-change | oper-status-change | neighbor
          discovery | control-channel-status}
```

3.9.2 Command Descriptions

3.9.2.1 VPRN Show Commands



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

egress-label

Syntax `egress-label start-label [end-label]`

Context `show>service`

Description This command displays services using the range of egress labels. If only the mandatory start-label parameter is specified, only services using the specified label are displayed.

If both start-label and end-label parameters are specified, the services using the range of labels X where start-label <= X <= end-label are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *start-label* — The starting egress label value for which to display services using the label range. If only start-label is specified, services only using start-label are displayed.

Values 0,18432 to 131071

end-label — The ending egress label value for which to display services using the label range.

Values 18432 to 131071

Output The following output is an example of show egress label information, and [Table 48](#) 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 48 Show Service Egress Field Descriptions

Label	Description
Svc Id	The ID that identifies a service.
Sdp Id	The ID that identifies an SDP.
Type	Indicates whether the SDP binding is a spoke or a mesh.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
E. Lbl	The VC label used by this device to send packets to the far-end device in this service by the SDP.
Number of bindings found	The total number of SDP bindings that exist within the specified egress label range.

ingress-label

Syntax **ingress-label** *start-label* [*end-label*]

Context show>service

Description Display services using the range of ingress labels.

If only the mandatory *start-label* parameter is specified, only services using the specified label are displayed.

If both *start-label* and *end-label* parameters are specified, the services using the range of labels X where *start-label* <= X <= *end-label* are displayed.

Use the **show router vprn-service-id ldp bindings** command to display dynamic labels.

Parameters *start-label* — The starting ingress label value for which to display services using the label range. If only *start-label* is specified, services only using *start-label* are displayed.

Values 0, 2048 to 131071

end-label — 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 show service ingress label information, and [Table 49](#) 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 49 Show Service Ingress-Label Field Descriptions

Label	Description
Svc ID	The service identifier.
SDP Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
I.Lbl	The ingress label used by the far-end device to send packets to this device in this service by the SDP.
E.Lbl	The egress label used by this device to send packets to the far-end device in this service by the SDP.
Number of bindings found	The number of SDP bindings within the label range specified.

sap-using

Syntax **sap-using** [msap] [dyn-script] [description]

```

sap-using [sap sap-id] [vlan-translation | anti-spoof] [description]
sap-using {ingress | egress} atm-td-profile td-profile-id
sap-using {ingress | egress} filter any-filter-id
sap-using {ingress | egress} qos-policy qos-policy-id [msap]
sap-using etree

```

Context show>service

Description This command displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

Parameters **msap** — Displays MSAPs associated to the service. This parameter applies to the 7450 ESS or 7750 SR only.

dyn-script — Displays dynamic service SAPs information.

description — Displays SAP description.

sap-id — Specifies the physical port identifier portion of the SAP definition.

vlan-translation — Displays VLAN translation information.

anti-spoof — Displays anti-spoof information.

ingress — Specifies matching an ingress policy.

egress — Specifies matching an egress policy.

td-profile-id — Displays SAPs using this traffic description.

Values 1 to 1000

any-filter-id — The ingress or egress filter policy ID for which to display matching SAPs. 64 characters maximum.

Values 1 to 65535

qos-policy-id — The ingress or egress QoS Policy ID for which to display matching SAPs. 64 characters maximum.

Values 1 to 65535

etree — Specifies matching of SAPs configured as E-Tree SAPs and the corresponding role in the E-Tree services: Leaf-AC, Root-AC or Root-leaf-tag SAPs. SAPs listed as Root-leaf-tag Disabled and Leaf-Ac Disabled function as Root-AC SAPs.

Output The following output is an example of service SAP information, and [Table 50](#) 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#

*A:ALA-12# show service sap-using egress atm-td-profile 2

Service Access Point Using ATM Traffic Profile 2

```

=====
PortId SvcId I.QoS I.Fltr E.QoS E.Fltr A.Pol Adm Opr
-----

```

```

5/1/1:0/11 511111 2 none 2 none none Up Up
5/1/1:0/12 511112 2 none 2 none none Up Up
5/1/1:0/13 511113 2 none 2 none none Up Up
5/1/1:0/14 511114 2 none 2 none none Up Up
5/1/1:0/15 511115 2 none 2 none none Up Up
5/1/1:0/16 511116 2 none 2 none none Up Up
5/1/1:0/17 511117 2 none 2 none none Up Up
5/1/1:0/18 511118 2 none 2 none none Up Up
5/1/1:0/19 511119 2 none 2 none none Up Up
5/1/1:0/20 511120 2 none 2 none none Up Up
5/1/1:0/21 511121 2 none 2 none none Up Up
5/1/1:0/22 511122 2 none 2 none none Up Up
5/1/1:0/23 511123 2 none 2 none none Up Up
5/1/1:0/24 511124 2 none 2 none none Up Up
5/1/1:0/25 511125 2 none 2 none none Up Up ...
-----

```

*A:ALA-12#

Table 50 Show Service SAP Field Descriptions

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.

Table 50 Show Service SAP Field Descriptions (Continued)

Label	Description (Continued)
Opr	The actual state of the SAP.

sdp-using

Syntax	sdp-using [<i>sdp-id[:vc-id]</i>] far-end <i>ip-address</i>]
Context	show>service
Description	Display services using SDP or far-end address options.
Parameters	<i>sdp-id</i> — Displays only services bound to the specified SDP ID. Values 1 to 17407 <i>vc-id</i> — The virtual circuit identifier. Values 1 to 4294967295 <i>ip-address</i> — Displays only services matching with the specified far-end IP address. 64 characters maximum. Default Services with any far-end IP address.
Output	The following output is an example of SDP-using information, and Table 51 describes the output fields.

Sample Output

```
*A:ALA-1# show service sdp-using 300
=====
Service Destination Point (Sdp Id : 300)
=====
SvcId      SdpId      Type Far End      Opr State I.Label  E.Label
-----
1          300:1      Mesh 10.0.0.13    Up      131071 131071
2          300:2      Spok 10.0.0.13    Up      131070 131070
100        300:100    Mesh 10.0.0.13    Up      131069 131069
101        300:101    Mesh 10.0.0.13    Up      131068 131068
102        300:102    Mesh 10.0.0.13    Up      131067 131067
-----
Number of SDPs : 5
-----
*A:ALA-1#

A:ALA-48# show service sdp-using
=====
SDP Using
=====
SvcId      SdpId      Type Far End      Opr State I.Label  E.Label
-----
```

3	2:3	Spok	10.20.1.2	Up	n/a	n/a
103	3:103	Spok	10.20.1.3	Up	131067	131068
103	4:103	Spok	10.20.1.2	Up	131065	131069
105	3:105	Spok	10.20.1.3	Up	131066	131067

Number of SDPs : 4

A:ALA-48

Table 51 Show Service SDP-Using Field Descriptions

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] [mirror] [apipe] [fpipe] [ipipe] [cpipe] [etree] [vsd] [b-vpls] [i-vpls] [m-vpls] [sdp sdp-id] [customer customer-id] [origin creation-origin]`
- Context** show>service
- Description** 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 (applies only to the 7450 ESS and 7750 SR).
 - ies** — Displays matching IES instances.
 - vpls** — Displays matching VPLS instances (applies only to the 7450 ESS and 7750 SR).
 - vprn** — Displays matching VPRN services.
 - mirror** — Displays mirror services.
 - apipe** — Displays matching Apipe services. (applies only to the 7450 ESS and 7750 SR).
 - fpipe** — Displays matching Fpipe services (applies only to the 7450 ESS and 7750 SR).

ipipe — Displays matching Ipipe services (applies only to the 7450 ESS and 7750 SR).

cpipe — Displays matching Cpipe services.

etree — Displays etree service information.

vsd — Displays matching VSD services.

b-vpls — Displays matching B-VPLS services.

i-vpls — Displays matching I-VPLS services.

m-vpls — Displays matching M-VPLS services.

sdp-id — Displays only services bound to the specified SDP ID.

Values 1 to 17407

Default Services bound to any SDP ID.

customer-id — Displays services only associated with the specified customer ID.

Values 1 to 2147483647

Default Services associated with a customer.

creation-origin — Specifies the method by which the service was created.

Values manual, multi-segment-p-w, dyn-script, vsd

Output The following output is an example of service-using information, and [Table 52](#) 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:del4# 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:del4#

```

Table 52 Show Service Service-Using Field Descriptions

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`

Context `show>service`

Description This command displays information for a particular service-id.

Parameters *service-id* — The unique service identification number that identifies the service in the service domain.

Values 1 to 2148278316, *svc-name*: 64 chars max

all

Syntax	all
Context	show>service>id
Description	This command displays detailed information for all aspects of the service.
Output	The following output is an example of show all service-ID information, and Table 53 describes the output fields.

Sample Output



Note: Ing ipv6 Fltr and Egr ipv6 Fltr are for the 7750 SR only.

```
A:PE# show service id 9 all
```

```
=====
Service Detailed Information
=====
Service Id       : 9                Vpn Id           : 0
Service Type     : VPRN
Name             : 9
Description      : (Not Specified)
Customer Id      : 1                Creation Origin   : manual
Last Status Change: 08/31/2018 16:07:55
Last Mgmt Change : 08/31/2018 16:07:55
Admin State      : Up               Oper State        : Up

Router Oper State : Up
Route Dist.       : 65536:9         VPRN Type         : regular
Oper Route Dist   : 65536:9
Oper RD Type      : configured
AS Number         : None            Router Id         : 1.1.1.1
ECMP              : Enabled         ECMP Max Routes   : 1
Max IPv4 Routes   : No Limit

Auto Bind Tunnel
Resolution        : disabled
Weighted ECMP     : Disabled        ECMP Max Routes   : 1

Max IPv6 Routes   : No Limit
Ignore NH Metric  : Disabled
Hash Label        : Disabled
Entropy Label     : Disabled
Vrf Target        : None
Vrf Import        : None
Vrf Export        : None
MVPN Vrf Target   : None
MVPN Vrf Import   : None
MVPN Vrf Export   : None
Car. Sup C-VPN    : Disabled
```

```
Label mode      : vrf
BGP VPN Backup  : Disabled
BGP Export Inactv : Disabled
```

```
SAP Count      : 1                SDP Bind Count    : 0
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)
Ingress uRPF check : enabled

```

Vxlan Tunnel Termination

Tunnel	Term	IP	FPE	ID	Last Change

No Matching Entries

Service Destination Points (SDPs)

No Matching Entries

Service Access Points

SAP 1/1/10:1

Service Id	: 9		
SAP	: 1/1/10:1	Encap	: q-tag
Description	: (Not Specified)		
Admin State	: Up	Oper State	: Up
Flags	: None		
Multi Svc Site	: None		
Last Status Change	: 08/31/2018 16:07:55		
Last Mgmt Change	: 08/31/2018 16:07:55		
Sub Type	: regular		
Dot1Q Ethertype	: 0x8100	QinQ Ethertype	: 0x8100
Split Horizon Group:	(Not Specified)		
Admin MTU	: 1518	Oper MTU	: 1518
Ingr IP Fltr-Id	: n/a	Egr IP Fltr-Id	: n/a
Ingr Mac Fltr-Id	: n/a	Egr Mac Fltr-Id	: n/a
Ingr IPv6 Fltr-Id	: n/a	Egr IPv6 Fltr-Id	: n/a
qinq-pbit-marking	: both		
Egr Agg Rate Limit	: max		
Q Frame-Based Acct	: Disabled	Limit Unused BW	: Disabled

Qinq-vlan-translation	: None	Qinq-vlan-translation Ids	: None
Acct. Pol	: None	Collect Stats	: Disabled
Anti Spoofing	: None	Dynamic Hosts	: Enabled
Avl Static Hosts	: 0	Tot Static Hosts	: 0
Calling-Station-Id	: n/a		

Application Profile: None
Transit Policy : None
AARP Id : None

Oper Group	: (none)	Monitor Oper Grp	: (none)
Host Lockout Plcy	: n/a		
Lag Link Map Prof	: (none)		
Bandwidth	: Not-Applicable		
Oper DCpu Prot Pol*	_default-access-policy		

ETH-CFM SAP specifics

Tunnel Faults	: n/a	AIS	: Disabled
MC Prop-Hold-Timer	: n/a		
Squelch Levels	: None		
Collect Lmm Stats	: Disabled		
LMM FC Stats	: None		
LMM FC In Prof	: None		

QOS

Ingress qos-policy	: 1	Egress qos-policy	: 1
Ingress FP QGrp	: (none)	Egress Port QGrp	: (none)
Ing FP QGrp Inst	: (none)	Egr Port QGrp Inst	: (none)
Shared Q plcy	: n/a	Multipoint shared	: Disabled
I. Sched Pol	: (Not Specified)		
E. Sched Pol	: (Not Specified)		
I. Policer Ctl Pol	: (Not Specified)		
E. Policer Ctl Pol	: (Not Specified)		
E. HS Sec. Shaper	: (Not Specified)		
I. QGrp Redir. List	: (Not Specified)		
E. QGrp Redir. List	: (Not Specified)		

Sap Aggregate Stats

	Packets	Octets
Ingress		
Aggregate Offered	: 0	0
Aggregate Forwarded	: 0	0
Aggregate Dropped	: 0	0
Egress		
Aggregate Forwarded	: 0	0
Aggregate Dropped	: 0	0

Sap Statistics

Last Cleared Time : 08/31/2018 16:29:08

	Packets	Octets
CPM Ingress	: 0	0

Forwarding Engine Stats

Dropped	: 0	0
Received Valid	: 0	0
Off. HiPrio	: 0	0
Off. LowPrio	: 0	0
Off. Uncolor	: 0	0
Off. Managed	: 0	0

Queueing Stats(Ingress QoS Policy 1)

Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

Queueing Stats(Egress QoS Policy 1)

Dro. In/InplusProf	: 0	0
Dro. Out/ExcProf	: 0	0
For. In/InplusProf	: 0	0
For. Out/ExcProf	: 0	0

Sap per Queue stats

	Packets	Octets

Ingress Queue 1 (Unicast) (Priority)

Off. HiPrio	: 0	0
Off. LowPrio	: 0	0
Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

Egress Queue 1

For. In/InplusProf	: 0	0
For. Out/ExcProf	: 0	0
Dro. In/InplusProf	: 0	0
Dro. Out/ExcProf	: 0	0

* indicates that the corresponding row element may have been truncated.

Service Interfaces

Interface

If Name	: int		
Admin State	: Up	Oper (v4/v6)	: Up/Down
Down Reason V6	: ifProtoOperDown		
Protocols	: None		
IP Addr/mask	: 20.0.0.2/30	Address Type	: Primary
IGP Inhibit	: Disabled	Broadcast Address	: Host-ones
HoldUp-Time	: 0	Track Srrp Inst	: 0
Description	: N/A		
Ignore Port State	: None		

Details

Description	: (Not Specified)		
If Index	: 3	Virt. If Index	: 3
Last Oper Chg	: 08/31/2018 16:07:55	Global If Index	: 257
Mon Oper Grp	: None		
Srrp En Rtnng	: Disabled	Hold time	: N/A
SAP Id	: 1/1/10:1		
TOS Marking	: Trusted	If Type	: VPRN
SNTP B.Cast	: False		
MAC Address	: 00:00:00:00:00:01	Mac Accounting	: Disabled
Ingress stats	: Disabled	IPv6 DAD	: Enabled
TCP MSS V4	: 0	TCP MSS V6	: 0
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
ARP Learn Unsolic*	: Disabled	ND Learn Unsolic*	: None
ARP Proactive Refr*	: Disabled	ND Proactive Refr*	: None
ARP Populate Host	: Disabled	ND Populate Host	: Disabled
ARP Route Tag	: 0	ND Route tag	: 0
IP MTU	: (default)		
IP Oper MTU	: 1500		
ARP Populate	: Disabled	Host Conn Verify	: Disabled
SHCV pol IPv4	: None		
Cflowd (unicast)	: None	Cflowd (multicast)	: None
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		
Rx Pkts	: 9	Rx Bytes	: 0
Rx V4 Pkts	: N/A	Rx V4 Bytes	: N/A
Rx V6 Pkts	: N/A	Rx V6 Bytes	: N/A
Tx Pkts	: 10	Tx Bytes	: 1060
Tx V4 Pkts	: 10	Tx V4 Bytes	: 1060
Tx V4 Discard Pkts	: 0	Tx V4 Discard Byt*	: 0
Tx V6 Pkts	: 0	Tx V6 Bytes	: 0
Tx V6 Discard Pkts	: 0	Tx V6 Discard Byt*	: 0
Tx DBcast Dis. Pk*	: 0	Tx DBcast Dis. By*	: 0
Mpls Rx Pkts	: 0	Mpls Rx Bytes	: 0
Mpls Tx Pkts	: 0	Mpls Tx Bytes	: 0
OperDCpuProtPlcy	: N/A		
Proxy ARP Details			
Rem Proxy ARP	: Disabled	Local Proxy ARP	: Disabled
Policies	: none		
Proxy Neighbor Discovery Details			
Local Pxy ND	: Disabled		
Policies	: none		
DHCP no local server			

DHCP Details

Description : (Not Specified)
Admin State : Down Lease Populate : 0
Gi-Addr : 20.0.0.2* Gi-Addr as Src Ip : Disabled
* = inferred gi-address from interface IP address

Action : Keep Trusted : Disabled

DHCP Proxy Details

Admin State : Down
Lease Time : N/A
Emul. Server : Not configured
DHCP Relay Proxy Details
Relay ucast : none
Siaddr ovr : N/A

Subscriber Authentication Details

Auth Policy : None

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

ISA Tunnel redundant next-hop information

Static Next-Hop :
Dynamic Next-Hop :

ICMP Details

Redirects	: Number - 100	Time (seconds)	- 10
Unreachables	: Number - 100	Time (seconds)	- 10
TTL Expired	: Number - 100	Time (seconds)	- 10
Parameter Problem	: Number - 100	Time (seconds)	- 10
ICMP Mask Reply	: True		

ICMPv6 Details

Packet Too Big	: Number - 100	Time (seconds)	- 10
Parameter Problem	: Number - 100	Time (seconds)	- 10
Redirects	: Number - 100	Time (seconds)	- 10
Time Exceeded	: Number - 100	Time (seconds)	- 10
Unreachables	: Number - 100	Time (seconds)	- 10

IPCP Address Extension Details

Peer IP Addr : Not configured
Peer Pri DNS Addr : Not configured
Peer Sec DNS Addr : Not configured

Admin Groups

No Matching Entries

```

-----
Srlg Groups
-----
No Matching Entries
-----
QoS Queue-Group Redirection Details
-----
Ingress FP QGrp   : (none)           Egress Port QGrp   : (none)
Ing FP QGrp Inst  : (none)           Egr Port QGrp Inst: (none)
=====
* indicates that the corresponding row element may have been truncated.
A:PE#

```

Sample Output

The following is a part of a sample output relevant to PW SAPs:

```

*A:Dut-B# show service id 3 all
...
-----
SAP pw-3:3
-----
Service Id       : 3
SAP              : pw-3:3                Encap           : q-tag
Description      : (Not Specified)
Admin State      : Up                    Oper State      : Up
Flags            : None
Multi Svc Site   : None
Last Status Change : 02/03/2015 18:04:39
Last Mgmt Change  : 02/03/2015 18:04:13
Sub Type         : regular
Split Horizon Group: (Not Specified)
Admin MTU        : 1518                  Oper MTU        : 1518
Ingr IP Fltr-Id  : n/a                  Egr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : n/a                  Egr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a                 Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : both
Egr Agg Rate Limit: max

Endpoint         : N/A

Vlan-translation : None
Acct. Pol        : None
Application Profile: None
Transit Policy   : None
Oper Group       : (none)                Monitor Oper Grp : (none)
Host Lockout Plcy : n/a
Ignore Oper Down  : Disabled
Lag Link Map Prof : (none)
Cflowd           : Disabled

Limit Unused BW   : Disabled
Collect Stats     : Disabled

```

Table 53 **Show All Service-ID Field Descriptions**

Label	Description
Service Detailed Information	
Service Id	The service identifier.
VPN Id	The number which identifies the VPN.
Customer Id	The customer identifier.
Last Status Change	The date and time of the most recent change in the administrative or operating status of the service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Admin State	The current administrative state.
Oper State	The current operational state.
Route Dist.	Displays the route distribution number.
AS Number	Displays the autonomous system number.
Router Id	Displays the router ID for this service.
ECMP	Displays equal cost multipath information.
ECMP Max Routes	Displays the maximum number of routes that can be received from the neighbors in the group or for the specific neighbor.
Max Routes	Displays the maximum number of routes that can be used for path sharing.
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.
SDP Id	The SDP identifier.
Description	Generic information about the service.
SAP Count	The number of SAPs specified for this service.
SDP Bind Count	The number of SDPs bound to this service.
Split Horizon Group	Name of the split horizon group for this service.
Description	Description of the split horizon group.
Last Changed	The date and time of the most recent management-initiated change to this split horizon group.

Table 53 Show All Service-ID Field Descriptions (Continued)

Label	Description
ETH-CFM Service Specifics	
Tunnel Faults	Whether tunnel faults are ignored or accepted.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.
Type	Indicates whether this Service SDP binding is a spoke or a mesh.
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the keepalive protocol.
Oper State	The current status of the keepalive protocol.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.

Table 53 Show All Service-ID Field Descriptions (Continued)

Label	Description
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. If the SDP type is GRE, then the following message displays: SDP delivery mechanism is not MPLS.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Number of SDPs	The total number SDPs applied to this service ID.
Service Access Points	
Service Id	The service identifier.
Port Id	The ID of the access port where this SAP is defined.
Description	Generic information about the SAP.
Encap Value	The value of the label used to identify this SAP on the access port.
Admin State	The desired state of the SAP.
Oper State	The operating state of the SAP.
Last Changed	The date and time of the last change.
Admin MTU	The 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.
SAP Statistics	

Table 53 Show All Service-ID Field Descriptions (Continued)

Label	Description
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.
Queueing Stats	
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.
SAP per Queue stats	
Ingress Queue 1	The index of the ingress QoS queue of this SAP.
High priority offered	The packets or octets count of the high priority traffic for the SAP.
High priority dropped	The number of high priority traffic packets/octets dropped.
Low priority offered	The packets or octets count of the low priority traffic.
Low priority dropped	The number of low priority traffic packets/octets dropped.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
Out profile forwarded	The number of out-of-profile octets (rate above CIR) forwarded.
Egress Queue 1	The index of the egress QoS queue of the SAP.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
In profile dropped	The number of in-profile packets or octets dropped for the SAP.
Out profile forwarded	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Out profile dropped	The number of out-of-profile packets or octets discarded.
State	Specifies whether DHCP relay is enabled on this SAP.
Info Option	Specifies whether Option 82 processing is enabled on this SAP.
Action	Specifies the Option 82 processing on this SAP or interface: keep, replace or drop.

Table 53 **Show All Service-ID Field Descriptions (Continued)**

Label	Description
Circuit ID	Specifies whether the If index is inserted in circuit ID sub-option of Option 82.
Remote ID	Specifies whether the far-end MAC address is inserted in Remote ID sub-option of Option 82.
Service Access Points	
Managed by Service	Specifies the service-id of the management VPLS managing 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.
Spoke SDPs	
Managed by Service	Specifies the service-id of the management VPLS managing this spoke SDP.
Managed by Spoke	Specifies the sap-id inside the management VPLS managing this spoke SDP.
Prune state	Specifies the STP state inherited from the management VPLS.
Peer Pw Bits	Indicates the bits set by the LDP peer when there is a fault on its side of the pseudowire. LAC failures occur on the SAP that has been configured on the pipe service, PSN bits are set by SDP-binding failures on the pipe service. The pwNotForwarding bit is set when none of the above failures apply, such as an MTU mismatch failure. This value is only applicable if the peer is using the pseudowire status signaling method to indicate faults. 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.
IPCP Address Extension Details	
Peer IP Addr	Specifies the remote IP address to be assigned to the far-end of the associated PPP/MLPPP link via IPCP extensions.
Peer Pri DNS Addr	Specifies a unicast IPv4 address for the primary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions.
Peer Sec DNS Addr	Specifies a unicast IPv4 address for the secondary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions.

authentication

Syntax **authentication**

Context show>service>id

Description This command enables the context to display subscriber authentication information.

statistics

Syntax statistics [*policy name*] [*sap sap-id*]

Context show>service>id>authentication

Description This command displays session authentication statistics for this service.

Parameters *name* — Specifies the subscriber authentication policy statistics to display. 32 characters maximum.

sap-id — Specifies the SAP ID statistics to display.

Output The following output is an example of show service ID authentication information.

Sample Output

```
*A:ALA-1# show service id 11 authentication statistics
=====
Authentication statistics
=====
Interface / SAP                Authentication    Authentication
                               Successful        Failed
-----
abc-11-10.1.0.254             1582            3
-----
Number of entries: 1
=====
*A:ALA-1#
```

arp

Syntax arp [*ip-address*] | [*mac ieee-address*] | [*sap sap-id*] | [*interface ip-int-name*] | [*sdp sdp-id:vc-id*] [*summary*]

Context show>service>id

Description This command displays the ARP table for the IES instance.

Parameters *ip-address* — Displays only ARP entries in the ARP table with the specified IP address.

Values a.b.c.d

Default All IP addresses.

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.

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

Default All MAC addresses.

sap-id — Displays SAP information for the specified SAP ID.

sdp-id:vc-id — Specifies the SDP ID and VC ID.

Values sdp-id: 1 to 17407
vc-id: 1 to 4294967295

ip-int-name — The IP interface name for which to display matching ARPs. 32 characters maximum.

Output The following output is an example of show service ID information, and [Table 54](#) describes the output fields.

Sample Output

```
*A:ALA-12# show service id 2 arp
=====
ARP Table
=====
IP Address      MAC Address      Type    Age      Interface      Port
-----
10.11.1.1       00:03:fa:00:08:22 Other    00:00:00 ies-100-10.11.1 1/1/11:0
=====
*A:ALA-12#
```

Table 54 **Show Service-ID ARP Field Descriptions**

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.

arp-host

Syntax	arp-host [wholesaler <i>service-id</i>] [sap <i>sap-id</i> interface <i>interface-name</i> ip-address <i>ip-address</i>[/<i>mask</i>] mac <i>ieee-address</i> {[port <i>port-id</i>] [no-inter-dest-id inter-dest-id <i>inter-dest-id</i>]}] [detail] arp-host statistics [sap <i>sap-id</i> interface <i>interface-name</i>] arp-host summary [interface <i>interface-name</i> saps]
Context	show>service>id
Description	This command displays ARP host related information.
Parameters	<i>service-id</i> — The VPRN service ID of the wholesaler. When specified in this context, SAP, SDP, interface, IP address and MAC parameters are ignored (applies only to the7750 SR). Values 1 to 2148007980, <i>svc-name</i>: 64 chars max <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. <i>interface-name</i> — Displays information for the specified IP interface. 32 characters maximum. <i>ip-address</i>[/<i>mask</i>] — Displays information associated with the specified IP address. Values ip-address: a.b.c.d. mask: 1 to 32 <i>ieee-address</i> — Specifies the MAC address. Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx <i>port-id</i> — Specifies the port ID. Values slot, mda, port no-inter-dest-id — Displays the information about no intermediate destination ID. <i>inter-dest-id</i> — Indicates the intermediate destination identifier received from either the DHCP or the RADIUS server or the local user database. 32 characters maximum detail — Displays detailed information. statistics — Displays ARP host statistics. summary — Displays summary information. saps — Displays SAP ARP host related information.
Output	The following output is an example of arp-host command information.

Sample Output

```
*A:Dut-C# show service id 2 arp-host
=====
ARP host table, service 2
=====
```

IP Address	Mac Address	Sap Id	Remaining Time	MC Stdbby
10.128.1.2	00:80:00:00:00:01	2/1/5:2	00h04m41s	
10.128.1.3	00:80:00:00:00:02	2/1/5:2	00h04m42s	
10.128.1.4	00:80:00:00:00:03	2/1/5:2	00h04m43s	
10.128.1.5	00:80:00:00:00:04	2/1/5:2	00h04m44s	
10.128.1.6	00:80:00:00:00:05	2/1/5:2	00h04m45s	
10.128.1.7	00:80:00:00:00:06	2/1/5:2	00h04m46s	
10.128.1.8	00:80:00:00:00:07	2/1/5:2	00h04m47s	
10.128.1.9	00:80:00:00:00:08	2/1/5:2	00h04m48s	
10.128.1.10	00:80:00:00:00:09	2/1/5:2	00h04m49s	
10.128.1.11	00:80:00:00:00:0a	2/1/5:2	00h04m50s	

Number of ARP hosts : 10

*A:Dut-C#

*A:Dut-C# show service id 2 arp-host ip-address 10.128.1.2 detail

ARP hosts for service 2

Service ID : 2
IP Address : 10.128.1.2
MAC Address : 00:80:00:00:00:01
SAP : 2/1/5:2
Remaining Time : 00h04m58s

Sub-Ident : "alu_1_2"
Sub-Profile-String : ""
SLA-Profile-String : ""
App-Profile-String : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
RADIUS-User-Name : "10.128.1.2"

Session Timeout (s) : 301
Start Time : 02/09/2009 16:35:07
Last Auth : 02/09/2009 16:36:34
Last Refresh : 02/09/2009 16:36:38
Persistence Key : N/A

Number of ARP hosts : 1

*A:Dut-C#

*A:Dut-C# show service id 2 arp-host statistics

ARP host statistics

Num Active Hosts : 20
Received Triggers : 70
Ignored Triggers : 10
Ignored Triggers (overload) : 0
SHCV Checks Forced : 0
Hosts Created : 20
Hosts Updated : 40

```

Hosts Deleted           : 0
Authentication Requests Sent : 40
=====
*A:Dut-C#

*A:Dut-C# show service id 2 arp-host summary
=====
ARP host Summary, service 2
=====
Sap                Used        Provided    Admin State
-----
sap:2/1/5:2        20          8000       inService
-----
Number of SAPs : 1
-----
=====
*A:Dut-C#

```

base

Syntax	base [msap]
Context	show>service>id
Description	Displays basic information about the service ID including service type, description, SAPs and SDPs.
Parameters	msap — Keyword to display MSAPs.
Output	The following output is an example of show service ID base command information, and Table 55 describes the output fields.

Sample Output

```

*A:SetupCLI# show service id 3 base
=====
Service Basic Information
=====
Service Id           : 3                Vpn Id           : 0
Service Type         : VPRN
Name                  : (Not Specified)
Description           : (Not Specified)
Customer Id          : 1
Last Status Change   : 10/08/2009 04:55:01
Last Mgmt Change     : 10/08/2009 06:48:38
Admin State           : Down             Oper State        : Down

Route Dist.           : None              VPRN Type         : regular
AS Number             : None              Router Id          : 10.20.30.40
ECMP                   : Enabled            ECMP Max Routes    : 1
Max IPv4 Routes        : No Limit           Auto Bind           : MPLS
Max IPv6 Routes        : No Limit
Ignore NH Metric       : Disabled
Hash Label             : Enabled

```

```

Vrf Target      : None
Vrf Import      : None
Vrf Export      : None
MVPN Vrf Target : None
MVPN Vrf Import : None
MVPN Vrf Export : None

SAP Count       : 0                      SDP Bind Count : 1
-----
Service Access & Destination Points
-----
Identifier                               Type      AdmMTU  OprMTU  Adm  Opr
-----
sdp:2000:1 S(101.101.101.101)           TLDP      1500    1500    Up   Down
=====
*A:SetupCLI#

```

Table 55 **Show Service-ID Base Field Descriptions**

Label	Description
Service Id	The service identifier.
Vpn Id	Specifies the VPN ID assigned to the service.
Service Type	Specifies the type of service.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Adm	The desired state of the service.
Oper	The operating state of the service.
Mtu	The largest frame size (in octets) that the service can handle.
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.

Table 55 Show Service-ID Base Field Descriptions (Continued)

Label	Description
Opr	The operating state of the SDP.

dhcp

Syntax	dhcp
Context	show>service>id
Description	This command enables the context to display DHCP information for the specified service.

lease-state

Syntax	lease-state [wholesaler <i>service-id</i>] [sap <i>sap-id</i> sdp <i>sdp-id:vc-id</i> interface <i>interface-name</i> ip-address <i>ip-address[/mask]</i> chaddr <i>ieee-address</i> mac <i>ieee-address</i> {[port <i>port-id</i>] [no-inter-dest-id <i>inter-dest-id</i>]}] [session { none ipoe }] [detail]
Context	show>service>id>dhcp
Description	This command displays DHCP lease state related information.
Parameters	<i>service-id</i> — The VPRN service ID of the wholesaler. When specified in this context, SAP, SDP, interface, IP address and MAC parameters are ignored (applies only to the 7750 SR). Values 1 to 2148007980, <i>svc-name: 64 chars max</i> <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. <i>sdp-id</i> — The SDP identifier. Values 1 to 17407 <i>vc-id</i> — The virtual circuit ID on the SDP ID for which to display information. Values 1 to 4294967295 <i>interface-name</i> — Displays information for the specified IP interface. 32 characters maximum. <i>ip-address</i> [<i>/mask</i>] — Displays information associated with the specified IP address. Values a.b.c.d mask: 1 to 32 chaddr <i>ieee-address</i> — Specifies the MA address of the DHCP lease state. Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx (cannot be all zeros)

mac *ieee-address* — Specifies the MAC address of the DHCP lease-state.

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx (cannot be all zeros)

port-id — Specifies the port identifier.

inter-dest-id — Shows DHCPv4 lease states for hosts that are associated with an IPoE session or for hosts that are not associated with an IPoE session. 32 characters maximum.

session — Shows DHCPv4 lease states for hosts that are associated with an IPoE session or for hosts that are not associated with an IPoE session.

Values ipoe, none

detail — Displays detailed information.

Output See the following sections for show command output.

- [Lease-State Sample Output](#)
- [Routed CO Sample Output](#)
- [Wholesaler/Retailer Sample Output](#)

Lease-State Sample Output

```
*A:ALA-48>config# show service id 101 dhcp lease-state
=====
DHCP lease state table, service 101
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining      Lease      MC
                  LifeTime      Origin      Stdby
-----
10.1.1.52       00:00:1f:bd:00:bb lag-1:101       00h02m56s     DHCP-R
10.3.2.62       00:00:1f:bd:00:c6 lag-1:105       00h02m59s     Radius
-----
Number of lease states : 2
=====
*A:ALA-48>config#

*A:ALA-48>config# show service id 105 dhcp lease-state wholesaler 101
=====
DHCP lease state table, service 105
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining      Lease      MC
                  LifeTime      Origin      Stdby
-----
Wholesaler 101 Leases
-----
10.3.2.62       00:00:1f:bd:00:c6 lag-1:105       00h00m39s     Radius
-----
Number of lease states : 1
=====
*A:ALA-48>config#
```

Routed CO Sample Output


```
A:ALA-_Dut-A# show service id 13 dhcp lease-state
=====
DHCP lease state table, service 13
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining   Lease   MC
                  LifeTime      Origin      Stdbby
-----
10.13.40.1      00:00:00:00:00:13  1/1/1:13      00h00m58s   Radius
-----
Number of lease states : 1
=====
A:ALA-_Dut-A#

A:ALA-_Dut-A# show service id 13 dhcp lease-state detail
=====
DHCP lease states for service 13
=====
Service ID      : 13
IP Address      : 10.13.40.1
Mac Address     : 00:00:00:00:00:13
Subscriber-interface : ies-13-13.13.1.1
Group-interface : intf-13
SAP             : 1/1/1:13
Remaining Lifetime : 00h00m58s
Persistence Key  : N/A

Sub-Ident       : "TEST"
Sub-Profile-String : "ADSL GO"
SLA-Profile-String : "BE-Video"
Lease ANCP-String : ""

Sub-Ident origin : Radius
Strings origin   : Radius
Lease Info origin : Radius

Ip-Netmask      : 255.255.0.0
Broadcast-Ip-Addr : 10.13.255.255
Default-Router   : N/A
Primary-Dns      : 10.13.254.254
Secondary-Dns    : 10.13.254.253

ServerLeaseStart : 12/24/2006 23:48:23
ServerLastRenew  : 12/24/2006 23:48:23
ServerLeaseEnd    : 12/24/2006 23:49:23
Session-Timeout   : 0d 00:01:00
DHCP Server Addr  : N/A

Persistent Relay Agent Information
  Circuit Id      : ancstb6_Dut-A|13|intf-13|0|13
  Remote Id       : stringtest
-----
Number of lease states : 1
=====
A:ALA-_Dut-A#
```

Wholesaler/Retailer Sample Output

```
A:ALA-_Dut-A# show service id 2000 dhcp lease-state detail
```

```

=====
DHCP lease states for service 2000
=====
Wholesaler 1000 Leases
-----
Service ID           : 1000
IP Address           : 10.13.1.254
Mac Address          : 00:00:00:00:00:13
Subscriber-interface : whole-sub
Group-interface      : intf-13
Retailer             : 2000
Retailer If          : retail-sub
SAP                  : 1/1/1:13
Remaining Lifetime   : 00h09m59s
Persistence Key       : N/A

Sub-Ident             : "TEST"
Sub-Profile-String    : "ADSL GO"
SLA-Profile-String    : "BE-Video"
Lease ANCP-String     : ""

Sub-Ident origin      : Retail DHCP
Strings origin        : Retail DHCP
Lease Info origin     : Retail DHCP

Ip-Netmask            : 255.255.0.0
Broadcast-Ip-Addr     : 10.13.255.255
Default-Router        : N/A
Primary-Dns           : N/A
Secondary-Dns         : N/A

ServerLeaseStart      : 12/25/2006 00:29:41
ServerLastRenew       : 12/25/2006 00:29:41
ServerLeaseEnd        : 12/25/2006 00:39:41
Session-Timeout       : 0d 00:10:00
DHCP Server Addr      : 10.232.237.2

Persistent Relay Agent Information
  Circuit Id          : 1/1/1:13
  Remote Id           : stringtest
-----
Number of lease states : 1
=====
A:ALA-_Dut-A#

```

statistics

Syntax	statistics [[sap sap-id] [sdp sdp-id:vc-id] [interface interface-name]]
Context	show>service>id>dhcp
Description	This command displays DHCP statistics information.
Parameters	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition.

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-name — Displays information for the specified IP interface.

Output The following output is an example of DHCP statistics, and [Table 56](#) 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 56 Show DHCP Statistics Field Descriptions

Label	Description
Received Packets	The number of packets received from the DHCP clients.
Transmitted Packets	The number of packets transmitted to the DHCP clients.
Received Malformed Packets	The number of corrupted/invalid packets received from the DHCP clients.
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients. In this case, a frame is dropped due to the client sending a DHCP packet with Option 82 filled in before “trust” is set under the DHCP interface command.
ClientPacketsDiscarded	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.

Table 56 Show DHCP Statistics Field Descriptions (Continued)

Label	Description
ServerPacketsRelayed	The number of packets received from the DHCP server that were forwarded.
ServerPacketsSnooped	The number of packets received from the DHCP server that were snooped.

gsmp

Syntax	gsmp
Context	show>service>id
Description	This command displays GSMP information.

neighbors

Syntax	neighbors [<i>group name</i> [<i>ip-address</i>]]
Context	show>service>id>gsmp
Description	This command displays GSMP neighbor information.
Parameters	group — A GSMP group defines a set of GSMP neighbors which have the same properties. <i>name</i> — Specifies a GSMP group name is unique only within the scope of the service in which it is defined. 32 characters maximum. <i>ip-address</i> — Specifies the ip-address of the neighbor. Values a.b.c.d (unicast address only)
Output	These commands show the configured neighbors per service, regardless of the fact there exists an open TCP connection with this neighbor. The admin state is shown because for a neighbor to be admin enabled, the service, gsmp node, group node and the neighbor node in this service must all be in 'no shutdown' state. Session gives the number of sessions (open TCP connections) for each configured neighbor.

Output Sample

The following show an example of GSMP neighbors output.

```
A:active>show>service>id>gsmp# neighbors
=====
GSMP neighbors
=====
Group                               Neighbor                AdminState  Sessions
-----
dslam1                             192.168.1.2            Enabled     0
dslam1                             192.168.1.3            Enabled     0
```

```

-----
Number of neighbors shown: 2
=====
A:active>show>service>id>gsmp#

A:active>show>service>id>gsmp# neighbors group dslam1
=====
GSMP neighbors
=====
Group                Neighbor                AdminState  Sessions
-----
dslam1                192.168.1.2            Enabled     0
dslam1                192.168.1.3            Enabled     0
-----
Number of neighbors shown: 2
=====
A:active>show>service>id>gsmp#

A:active>show>service>id>gsmp# neighbors group dslam1 192.168.1.2
=====
GSMP neighbors
=====
Group                Neighbor                AdminState  Sessions
-----
dslam1                192.168.1.2            Enabled     0
=====
A:active>show>service>id>gsmp#

```

sessions

- Syntax** **sessions** {**group** *name* | **neighbor** *ip-address* **port** *port-number* [**association** | **statistics**]}
- Context** show>service>id>gsmp
- Description** This command displays GSMP sessions information.
- Parameters**
- group** — A GSMP group defines a set of GSMP neighbors which have the same properties.
 - name** — Specifies a GSMP group name is unique only within the scope of the service in which it is defined. 32 characters maximum.
 - ip-address** — Specifies the IP address of the neighbor.
 - Values** a.b.c.d (unicast address only)
 - port-number** — Specifies the neighbor TCP port number use for this ANCP session.
 - Values** 0 to 65535
 - association** — Displays to what object the ANCP-string is associated.
 - statistics** — Displays statistics information about an ANCP session known to the system.
- Output** The following output sample is an example of GSMP sessions information.

Sample Output

This show command gives information about the open TCP connections with DSLAMs.

```
A:active>show>service>id>gsmp# sessions
```

```
=====
```

```
GSMP sessions for service 999 (VPRN)
```

```
=====
```

```
Port      Ngbr-IPAddr      Gsmp-Group
```

```
-----
```

```
40590    192.168.1.2      dslam1
```

```
-----
```

```
Number of GSMP sessions : 1
```

```
=====
```

```
A:active>show>service>id>gsmp#
```

```
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590
```

```
=====
```

```
GSMP sessions for service 999 (VPRN), neighbor 192.168.1.2, Port 40590
```

```
=====
```

```
State           : Established
```

```
Peer Instance    : 1
```

```
Sender Instance  : a3cf58
```

```
Peer Port        : 0
```

```
Sender Port      : 0
```

```
Peer Name        : 12:12:12:12:12:12
```

```
Sender Name      : 00:00:00:00:00:00
```

```
Timeouts         : 0
```

```
Max. Timeouts    : 3
```

```
Peer Timer       : 100
```

```
Sender Timer      : 100
```

```
Capabilities      : DTD OAM
```

```
Conf Capabilities : DTD OAM
```

```
Priority Marking   : dscp nc2
```

```
Local Addr.       : 192.168.1.4
```

```
Conf Local Addr.  : N/A
```

```
=====
```

```
A:active>show>service>id>gsmp#
```

```
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 association
```

```
=====
```

```
ANCP-Strings
```

```
=====
```

```
ANCP-String                                           Assoc. State
```

```
-----
```

```
No ANCP-Strings found
```

```
=====
```

```
A:active>show>service>id>gsmp#
```

```
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 statistics
```

```
=====
```

```
GSMP session stats, service 999 (VPRN), neighbor 192.168.1.2, Port 40590
```

```
=====
```

```
Event           Received    Transmitted
```

```
-----
```

```
Dropped
```

```
0
```

```
0
```

```
Syn
```

```
1
```

```
1
```

```
Syn Ack
```

```
1
```

```
1
```

```
Ack
```

```
14
```

```
14
```

```
Rst Ack
```

```
0
```

```
0
```

```
Port Up
```

```
0
```

```
0
```

```
Port Down
```

```
0
```

```
0
```

```
OAM Loopback                                0          0
=====
A:active>show>service>id>gsmp#
```

The association command gives an overview of each ANCP string received from this session.

```
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 association
=====
ANCP-Strings
=====
ANCP-String                                     Assoc.
State
-----
7330-ISAM-E47 atm 1/1/01/01:19425.64048          ANCP    Up
-----
Number of ANCP-Strings : 1
=====
A:active>show>service>id>gsmp#
```

summary

- Syntax

summary [interface *interface-name* | saps]
- Context

show>service>id>dhcp
- Description

This command displays DHCP configuration summary information.
- Parameters

interface-name — Displays information for the specified IP interface.

Values 32 characters maximum.

saps — Displays SAPs per interface.
- Output

The following output is an example of DHCP summary information, and [Table 57](#) describes the output fields.

Sample Output

```
A:ALA-49# show service id 1 dhcp summary
=====
DHCP Summary, service 1
=====
Interface Name      Arp      Used/      Info      Admin
SapId/Sdp           Populate Provided Option    State
-----
SpokeSDP            No        0/0        Keep      Up
  sdp:spoke-3:4      No        0/0        Keep      Up
test                No        0/0        Keep      Up
  sap:1/1/4:50/5     No        0/0        Keep      Up
to-cel              No        0/0        Keep      Up
  sap:1/1/10:1       No        0/0        Keep      Up
-----
Interfaces: 3
=====
```

A:ALA-49#

Table 57 Show Service-ID DHCP Summary Field Descriptions

Label	Description
Sap/Sdp	The configuration identification, expressed by a string containing "card/mda/port/:logical-id".
Snoop	Yes — The packets received from the DHCP clients were snooped. No — The packets received from the DHCP clients were not snooped.
Used/Provided	Used — The number of lease-states that are currently in use on a specific interface, that is, the number of clients on that interface got an IP address by DHCP. This value is always less than or equal to the 'Provided' field. Provided — The lease-populate value that is configured for a specific interface.
Arp Reply Agent	Displays whether or not there is proper handling of received ARP requests from subscribers.
Info Option	Keep — The existing information is kept on the packet and the router does not add any additional information. Replace — On ingress, the existing information-option is replaced with the information-option from the router. Drop — The packet is dropped and an error is logged.
Admin State	Indicates the administrative state.

interface

Syntax	interface <i>[[ip-address ip-int-name] [interface-type] [detail] [family]] summary</i>
Context	show>service>id
Description	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	<i>ip-address</i> — The IP address of the interface for which to display information. Values 1.0.0.0 — 223.255.255.255 <i>ip-int-name</i> — The IP interface name for which to display information. 32 characters maximum. family — Specifies the family to display. Values ipv4, ipv6 interface-type — Specifies the interface type. Values subscriber, group, redundant

detail — Displays detailed IP interface information.

Default IP interface summary output

summary — Displays a summary.

Output The following output is an example of IP interface information, and [Table 58](#) describes the output fields.

Sample Output

```
A:cses-V96# /show service id 1 interface
- interface [{[<ip-address|ip-int-name>] [interface-type] [detail]
  [family]} | summary]
<ip-int-name|ip-ad*> : ip-int-name      - 32 chars max
                      ipv4-address     - a.b.c.d
                      ipv6-address     - x:x:x:x:x:x:x:x (eight 16-bit
                                      pieces)
                                      x:x:x:x:x:x.d.d.d.d
                                      x - [0..FFFF]H
                                      d - [0..255]D

<detail>              : keyword - adds details to the display
<family>              : ipv4|ipv6
<interface-type>      : subscriber|group|redundant
<summary>             : keyword - displays summary
```

Sample Output

```
*A:ALA-12# show service id 321 interface
=====
Interface Table
=====
Interface-Name          Type IP-Address      Adm   Opr   Type
-----
test                    Pri  10.11.1.1/24      Up    Up    IES
-----
Interfaces : 1
=====
*A:ALA-12#

A:ALA-49# show service id 88 interface detail
=====
Interface Table
=====
Interface
-----
If Name      : Sector A
Admin State  : Up
Protocols    : None
Oper State   : Down

IP Addr/mask : Not Assigned
-----
Details
-----
Description :
```

```

If Index      : 26
SAP Id        : 7/1/1.2.2
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
Lease Populate : 0
Trusted        : Disabled

ICMP Details
Redirects      : Number - 100
Unreachables   : Number - 100
TTL Expired    : Number - 100
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        : 10/1/2:0
TOS Marking   : Untrusted
SNTP B.Cast   : False
MAC Address   : Not configured.
IP MTU        : 1500
Arp Populate  : Disabled
Cflowd       : None

Virt. If Index : 27
If Type        : IES
IES ID         : 88
Arp Timeout    : 14400
ICMP Mask Reply : True

Proxy ARP Details
Proxy ARP      : Disabled
Local Proxy ARP : Disabled

DHCP Details
Admin State    : Up
Action         : Keep
Lease Populate : 0
Trusted        : 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#

```

```

*A:SetupCLI# show service id 3 interface "ab" detail
=====

```

```

Interface Table
=====
-----
Interface
-----
If Name      : ab
Admin State  : Up                               Oper (v4/v6)   : Down/--
Protocols    : None

IP Addr/mask : Not Assigned
-----
Details
-----
Description  : (Not Specified)
If Index     : 2                               Virt. If Index : 2
Last Oper Chg: 10/08/2009 07:07:58           Global If Index : 329
SDP Id       : spoke-2000:1

Spoke-SDP Details
Admin State  : Up                               Oper State    : Down
Hash Label   : Enabled
Flags        : SvcAdminDown SdpOperDown
               NoIngVCLLabel NoEgrVCLLabel

TOS Marking  : Trusted                         If Type       : VPRN
SNTP B.Cast  : False
MAC Address  : 76:6d:ff:00:00:00              Arp Timeout   : 14400
IP Oper MTU  : 0                             ICMP Mask Reply : True
Arp Populate : Disabled                       Host Conn Verify : Disabled
Cflowd      : None
LdpSyncTimer : None
LSR Load Bal*: system
uRPF Chk     : disabled
uRPF Fail By*: 0                             uRPF Chk Fail Pk*: 0

Proxy ARP Details
Rem Proxy ARP: Disabled                       Local Proxy ARP : Disabled
Policies     : none

Proxy Neighbor Discovery Details
Local Pxy ND : Disabled
Policies     : none

DHCP no local server

DHCP Details
Description  : (Not Specified)
Admin State  : Down                           Lease Populate : 0
Gi-Addr     : Not configured                 Gi-Addr as Src Ip: Disabled
Action      : Keep                           Trusted        : Disabled

DHCP Proxy Details
Admin State  : Down
Lease Time   : N/A
Emul. Server : Not configured

Subscriber Authentication Details
Auth Policy  : None

```

```

DHCP6 Relay Details
Description   : (Not Specified)
Admin State  : Down
Oper State   : Down
If-Id Option : None
Src Addr     : Not configured
Lease Populate : 0
Nbr Resolution : Disabled
Remote Id     : Disabled

DHCP6 Server Details
Admin State   : Down
Max. Lease States: 8000

ICMP Details
Redirects     : Number - 100
Unreachables  : Number - 100
TTL Expired   : Number - 100
Time (seconds) - 10
Time (seconds) - 10
Time (seconds) - 10

IPCP Address Extension Details
Peer IP Addr* : Not configured
Peer Pri DNS* : Not configured
Peer Sec DNS* : Not configured

Routed VPLS Details
VPLS Name     :
Binding Status : Up
-----
Interfaces : 1
=====
* indicates that the corresponding row element may have been truncated.
*A:SetupCLIP#

```

The Oper Hash Label and Hash Lbl Sig Cap spoke-sdp fields display when signal-capability is enabled and operational state of hash-label in datapath.

```

-----
Service Destination Points (SDPs)
-----
-----
Sdp Id 1:555 - (10.2.2.2)
-----
Description      : (Not Specified)
SDP Id           : 1:555
Spoke Descr      : (Not Specified)
VC Type          : Ether
Admin Path MTU   : 0
Far End          : 10.2.2.2
Tunnel Far End   : n/a
Hash Label       : Disabled
Oper Hash Label  : Disabled
Type             : Spoke
VC Tag           : n/a
Oper Path MTU    : 1568
Delivery         : MPLS
LSP Types        : RSVP
Hash Lbl Sig Cap : Disabled

Admin State      : Up
Acct. Pol        : None
Ingress Label    : 131065
Ingr Mac Fltr-Id : n/a
Ingr IP Fltr-Id  : n/a
Ingr IPv6 Fltr-Id : n/a
Admin ControlWord : Not Preferred
Admin BW(Kbps)   : 0
Last Status Change : 11/25/2010 13:06:14
Last Mgmt Change  : 11/24/2010 13:00:48
Oper State       : Up
Collect Stats    : Disabled
Egress Label     : 131059
Egr Mac Fltr-Id  : n/a
Egr IP Fltr-Id   : n/a
Egr IPv6 Fltr-Id : n/a
Oper ControlWord : False
Oper BW(Kbps)    : 0
Signaling        : TLDP
Force Vlan-Vc    : Disabled

```

```

Endpoint          : N/A                      Precedence       : 4
PW Status Sig     : Enabled
Class Fwding State : Down
Flags             : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : mplsRouterAlertLabel
Application Profile : None
Standby Sig Slave  : False

.....
.....
=====

```

Table 58 **Show Service-ID Interface Field Descriptions**

Label	Description
Interface-Name	The name used to refer to the interface.
Type	Specifies the interface type.
IP-Address	Specifies the IP address/IP subnet/broadcast address of the interface.
Adm	The desired state of the interface.
Opr	The operating state of the interface.
Interface	
If Name	The name used to refer to the interface.
Admin State	The desired state of the interface.
Oper State	The operating state of the interface.
IP Addr/mask	Specifies the IP address/IP subnet/broadcast address of the interface.
Ignore Port State	Indicates whether or not the tools perform service id service-id interface ip-int-name ignore-sap port-state command has been executed for a service interface directly connected to a SAP: none — The command has not been executed for or accepted by the interface. active — The command has been executed and accepted, and the port state check is currently being bypassed for the interface. pending — The command has been executed and accepted, but the port state for the interface is already operational.
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.

Table 58 Show Service-ID Interface Field Descriptions (Continued)

Label	Description (Continued)
Port Id	Specifies the SAP's port ID.
SNTP B.Cast	Specifies whether SNTP broadcast client mode is enabled or disabled.
Arp Timeout	Specifies the timeout for an ARP entry learned on the interface.
MAC Address	Specifies the 48-bit IEEE 802.3 MAC address.
ICMP Mask Reply	Specifies whether ICMP mask reply is enabled or disabled.
Cflowd	Specifies whether Cflowd collection and analysis on the interface 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.

ptp

Syntax	ptp
Context	show>service>id
Description	This command displays Precision Timing Protocol (PTP) information.

retailers

Syntax	retailers
Context	show>service>id
Description	This command displays the service ID of the retailer subscriber service to which this DHCP lease belongs.
Output	The following is sample output of the retailers command.

Sample Output

```
*A:ALA-48>config# show service id 101 retailers
=====
Retailers for service 101
=====
Retailer Svc ID                Num Static Hosts      Num Dynamic Hosts
```

```

-----
102                      3                      1
105                      0                      1
-----
Number of retailers : 2
=====
*A:ALA-48>config#

```

wholesalers

Syntax	wholesalers
Context	show>service>id
Description	This command displays service wholesaler information.
Output	The following is sample output of the wholesalers command.

Sample Output

```

*A:ALA-48>config# show service id 102 wholesalers
=====
Wholesalers for service 102
=====
Wholesaler Svc ID          Num Static Hosts          Num Dynamic Hosts
-----
101                      3                      1
-----
Number of wholesalers : 1
=====
*A:ALA-48>config#

```

Wholesaler information can also be displayed in the lease-state context.

```

*A:ALA-48>config# show service id 105 dhcp lease-state wholesaler 101
=====
DHCP lease state table, service 105
=====
IP Address      Mac Address      Sap/Sdp Id          Remaining   Lease   MC
                  LifeTime        Origin             Stdby
-----
Wholesaler 101 Leasesok
-----
10.3.2.62       00:00:1f:bd:00:c6 lag-1:105          00h00m39s   Radius
-----
Number of lease states : 1
=====
*A:ALA-48>config#

```

sap

Syntax	<pre> sap <i>sap-id</i> static-isids [<i>range-id range-id</i>] sap <i>sap-id</i> dist-cpu-protection [<i>detail</i>] sap <i>sap-id</i> detail sap <i>sap-id</i> encap-group <i>group-name</i> [<i>member encap-id</i>] [<i>encap-detail</i> <i>encap-stats</i>] sap <i>sap-id</i> encap-group sap <i>sap-id</i> [<i>atm</i> <i>base</i> <i>dhcp</i> <i>mc-ring</i> <i>mcac</i> <i>mrp</i> <i>qos</i> <i>sap-stats</i> <i>stats</i> <i>stp</i> <i>sub-</i> <i>mgmt</i> <i>ipsec-gw</i>] sap <i>sap-id</i> [<i>circuit-id circuit-id</i>] [<i>mac ieee-address</i>] [<i>remote-id remote-id</i>] host-lockout- policy [<i>summary</i>] sap <i>sap-id</i> queue-depth [<i>queue queue-id</i>] [<i>ingress</i> <i>egress</i>] sap sap <i>sap-id</i> static-isids mfib sap <i>sap-id</i> queue-group-redirect [<i>ingress</i> <i>egress</i>] </pre>
Context	show>service>id
Description	Displays information for the SAPs associated with the service. If no optional parameters are specified, a summary of all associated SAPs is displayed.
Parameters	<i>sap-id</i> — The ID that displays SAPs for the service. detail — Displays detailed information for the SAP. static-isids — Displays static ISIDs for the SAP. <i>range-id</i> — Specifies the range ID. Values 1 to 4294967295 dist-cpu-protection — Displays information about the distributed CPU protection policy and parameters associated with the SAP. <i>group-name</i> — Specifies the group name, up to 32 characters in length. <i>encap-id</i> — Specifies the encapsulation ID. Values 0 to 16777215 encap-detail — Displays encapsulation details for the specified encapsulation ID. encap-stats — Displays encapsulation statistics. encap-group — Displays the encapsulation group. atm — Displays ATM information. base — Displays base information. dhcp — Displays DHCP information. mc-ring — Displays MC ring information. mcac — Displays MCAC information.

mrp — Displays MRP information.

qos — Displays QoS information.

sap-stats — Displays SAP statistics.

stats — Displays statistics.

stp — Displays STP information.

sub-mgmt — Displays subscriber management information.

ipsec-gw — Displays IPSEC gateway information.

circuit-id — Specifies the circuit ID, up to 256 characters maximum.

ieee-address — Specifies the IEEE address, up to 30 characters maximum.

remote-id — Specifies the remote-ID, up to 256 characters maximum.

host-lockout-policy — Displays the host lockout policy.

summary — Displays summary information.

queue-id — Specifies the queue ID.

Values 1 to 32

ingress — Displays ingress information.

egress — Displays egress information.

mfib — Displays MFIB information.

queue-group-redirect — The output lists the queue group name and the instances configured in the related queue group redirect list. For each instance, the FP (for ingress) and port (for egress) is displayed. If there is a mismatch between the SAP and redirect list configuration and the queue group instance configuration, this is highlighted.

ingress — Displays information for the ingress policy.

egress — Displays information for the egress policy.

Output The following output is an example of SAP information, and [Table 59](#) describes the output fields.

Sample Output

```
A:PE# show service id 9 sap 1/1/10:1 detail

=====
Service Access Points (SAP)
=====
Service Id      : 9
SAP             : 1/1/10:1                      Encap           : q-tag
Description     : (Not Specified)
Admin State     : Up                            Oper State      : Up
Flags           : None
Multi Svc Site  : None
Last Status Change : 08/31/2018 16:07:55
```

```

Last Mgmt Change   : 08/31/2018 16:07:55
Sub Type           : regular
Dot1Q Ethertype    : 0x8100
Split Horizon Group: (Not Specified)
QinQ Ethertype     : 0x8100

Admin MTU          : 1518
Ingr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a
qinq-pbit-marking  : both
Egr Agg Rate Limit : max
Q Frame-Based Acct : Disabled
Qinq-vlan-translation : None
Acct. Pol          : None

Oper MTU           : 1518
Egr IP Fltr-Id    : n/a
Egr Mac Fltr-Id   : n/a
Egr IPv6 Fltr-Id  : n/a
Limit Unused BW    : Disabled
Qinq-vlan-translation Ids : None
Collect Stats      : Disabled

Anti Spoofing      : None
Avl Static Hosts   : 0
Calling-Station-Id : n/a
Dynamic Hosts      : Enabled
Tot Static Hosts   : 0

Application Profile: None
Transit Policy     : None
AARP Id            : None

Oper Group         : (none)
Host Lockout Plcy  : n/a
Lag Link Map Prof  : (none)
Bandwidth          : Not-Applicable
Oper Dcpu Prot Pol*: _default-access-policy

```

ETH-CFM SAP specifics

```

Tunnel Faults      : n/a
MC Prop-Hold-Timer : n/a
Squelch Levels     : None
Collect Lmm Stats  : Disabled
LMM FC Stats       : None
LMM FC In Prof     : None
AIS                : Disabled

```

QOS

```

Ingress qos-policy : 1
Ingress FP QGrp    : (none)
Ing FP QGrp Inst   : (none)
Shared Q plcy      : n/a
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
I. Policer Ctl Pol : (Not Specified)
E. Policer Ctl Pol : (Not Specified)
E. HS Sec. Shaper  : (Not Specified)
I. QGrp Redir. List: (Not Specified)
E. QGrp Redir. List: (Not Specified)
Egress qos-policy  : 1
Egress Port QGrp   : (none)
Egr Port QGrp Inst: (none)
Multipoint shared  : Disabled

```

Sap Aggregate Stats-----
Packets

Octets

```

Ingress
Aggregate Offered      : 0                0
Aggregate Forwarded    : 0                0
Aggregate Dropped      : 0                0

Egress
Aggregate Forwarded    : 0                0
Aggregate Dropped      : 0                0
-----
Sap Statistics
-----
Last Cleared Time      : 08/31/2018 16:29:08

                Packets                Octets
CPM Ingress           : 0                0

Forwarding Engine Stats
Dropped               : 0                0
Received Valid        : 0                0
Off. HiPrio           : 0                0
Off. LowPrio          : 0                0
Off. Uncolor          : 0                0
Off. Managed          : 0                0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                0
Dro. LowPrio          : 0                0
For. InProf           : 0                0
For. OutProf          : 0                0

Queueing Stats(Egress QoS Policy 1)
Dro. In/InplusProf    : 0                0
Dro. Out/ExcProf       : 0                0
For. In/InplusProf     : 0                0
For. Out/ExcProf       : 0                0
-----
Sap per Queue stats
-----
                Packets                Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                0
Off. LowPrio          : 0                0
Dro. HiPrio           : 0                0
Dro. LowPrio          : 0                0
For. InProf           : 0                0
For. OutProf          : 0                0

Egress Queue 1
For. In/InplusProf     : 0                0
For. Out/ExcProf       : 0                0
Dro. In/InplusProf     : 0                0
Dro. Out/ExcProf       : 0                0
=====
* indicates that the corresponding row element may have been truncated.
A:PE#

```

Table 59 **Show Service-ID SAP Field Descriptions**

Label	Description
Service Id	The service identifier.
SAP	The SAP and qtag.
Encap	The encapsulation type of the SAP.
Ethertype	Specifies an Ethernet type II Ethertype value.
Admin State	The administrative state of the SAP.
Oper State	The operating state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapPipeCelpAddr, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
Last Status Change	Specifies the time of the most recent operating status change to this SAP.
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SAP.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP.
Egress qos-policy	The egress QoS policy ID assigned to the SAP.
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP.
Egress Filter-Id	The egress filter policy ID assigned to the SAP.
Acct. Pol	The accounting policy ID assigned to the SAP.
Collect Stats	Specifies whether collect stats is enabled.
Dropped	The number of packets and octets dropped due to SAP state, ingress MAC or IP filter, same segment discard, bad checksum, etc.
Off. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Off. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.

Table 59 Show Service-ID SAP Field Descriptions (Continued)

Label	Description (Continued)
Off. Uncolor	The number of uncolored packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Dro. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to: MBS exceeded, buffer pool limit exceeded, etc.
Dro. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to: MBS exceeded, buffer pool limit exceeded, etc.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded by the ingress Qchip.
For. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
Dro. InProf	The number of in-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
Dro. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded by the egress Qchip.
For. OutProf	The number of out-of-profile packets and octets (rate above CIR) forwarded by the egress Qchip.
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 [:vc-id]] [detail]
 sdp far-end {ip-address | ipv6-address} [detail]
 sdp sdp-id [vc-id] l2tpv3
 sdp sdp-id [:vc-id] static-isids [range-id range-id]
 sdp sdp-id [:vc-id] static-isids mfib
 sdp sdp-id [:vc-id] [detail] vccv-bfd [session]
 sdp sdp-id [:vc-id] mrp

Context show>service>id

Description 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* [:*vc-id*] — Displays only information for the specified SDP ID.

Values sdp-id: 1 to 17407
vc-id: 1 to 4294967295

Default all SDPs

detail — Displays detailed SDP information.

ip-address | *ipv6-address* — Displays only SDPs matching with the specified far-end IP address. 64 characters maximum.

Default SDPs with any far-end IP address

l2tpv3 — Indicates that the user wants to display l2tpv3 specific information for SDPs that are of type l2tpv3.

static-isids — Specifies the I-Component service IDs created on the SDP.

range-id — Displays the service using the specified I-component Service ID (ISID).

Values 1 to 4294967295

mfib — Display MFIB related information. This parameter applies to the 7450 ESS or 7750 SR only.

vccv-bfd — Displays detailed information about the VCCV BFD session for a spoke SDP.

session — displays a summary of all VCCV sessions.

mrp — Displays detailed MRP information.

Output The following output is an example of SDP information, and [Table 60](#) 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
VC Type          : Ether
Admin Path MTU   : 0
Far End          : 10.20.1.2
Type             : Spoke
VC Tag           : n/a
Oper Path MTU    : 9186
Delivery         : MPLS

Admin State      : Up
Acct. Pol        : None
Ingress Label    : 2048
Ing mac Fltr     : n/a
Ing ip Fltr      : n/a
Oper State       : Up
Collect Stats    : Disabled
Egress Label     : 2048
Egr mac Fltr     : n/a
Egr ip Fltr      : n/a
```

Ingr ipv6 Fltr	: n/a	Egr ipv6 Fltr	: n/a
Admin ControlWord	: Not Preferred	Oper ControlWord	: False
Last Status Change	: 05/31/2007 00:45:43	Signaling	: None
Last Mgmt Change	: 05/31/2007 00:45:43		
Class Fwding State	: Up		
Flags	: None		
Peer Pw Bits	: None		
Peer Fault Ip	: None		
Peer Vccv CV Bits	: None		
Peer Vccv CC Bits	: None		
Max Nbr of MAC Addr	: No Limit	Total MAC Addr	: 0
Learned MAC Addr	: 0	Static MAC Addr	: 0
MAC Learning	: Enabled	Discard Unkwn Srce	: Disabled
MAC Aging	: Enabled		
L2PT Termination	: Disabled	BPDU Translation	: Disabled
MAC Pinning	: Disabled		
KeepAlive Information	:		
Admin State	: Disabled	Oper State	: Disabled
Hello Time	: 10	Hello Msg Len	: 0
Max Drop Count	: 3	Hold Down Time	: 10
Statistics	:		
I. Fwd. Pkts.	: 0	I. Dro. Pkts.	: 0
I. Fwd. Octs.	: 0	I. Dro. Octs.	: 0
E. Fwd. Pkts.	: 0	E. Fwd. Octets	: 0
MCAC Policy Name	:		
MCAC Max Unconst BW	: no limit	MCAC Max Mand BW	: no limit
MCAC In use Mand BW	: 0	MCAC Avail Mand BW	: unlimited
MCAC In use Opnl BW	: 0	MCAC Avail Opnl BW	: unlimited
Associated LSP LIST	:		
Lsp Name	: A_B_1		
Admin State	: Up	Oper State	: Up
Time Since Last Tr*	: 00h26m35s		
Lsp Name	: A_B_2		
Admin State	: Up	Oper State	: Up
Time Since Last Tr*	: 00h26m35s		
Lsp Name	: A_B_3		
Admin State	: Up	Oper State	: Up
Time Since Last Tr*	: 00h26m34s		
Lsp Name	: A_B_4		
Admin State	: Up	Oper State	: Up
Time Since Last Tr*	: 00h26m34s		
Lsp Name	: A_B_5		
Admin State	: Up	Oper State	: Up
Time Since Last Tr*	: 00h26m34s		
Lsp Name	: A_B_6		
Admin State	: Up	Oper State	: Up
Time Since Last Tr*	: 00h26m34s		
Lsp Name	: A_B_7		
Admin State	: Up	Oper State	: Up

Time Since Last Tr*: 00h26m34s

```

Lsp Name       : A_B_8
Admin State    : Up
Oper State     : Up
Time Since Last Tr*: 00h26m35s

```

```

Lsp Name       : A_B_9
Admin State    : Up
Oper State     : Up
Time Since Last Tr*: 00h26m34s

```

```

Lsp Name       : A_B_10
Admin State    : Up
Oper State     : Up
Time Since Last Tr*: 00h26m34s

```

Class-based forwarding :

```

Class forwarding : enabled
Default LSP      : A_B_10
Multicast LSP    : A_B_9

```

=====

FC Mapping Table

=====

FC Name	LSP Name
af	A_B_3
be	A_B_1
ef	A_B_6
h1	A_B_7
h2	A_B_5
l1	A_B_4
l2	A_B_2
nc	A_B_8

=====

Stp Service Destination Point specifics

```

Mac Move       : Blockable
Stp Admin State : Up
Core Connectivity : Down
Port Role      : N/A
Port Number    : 2049
Port Path Cost : 10
Admin Edge     : Disabled
Link Type      : Pt-pt
Root Guard     : Disabled
Last BPDU from : N/A
Designated Bridge : N/A

Stp Oper State : Down
Port State     : Forwarding
Port Priority   : 128
Auto Edge      : Enabled
Oper Edge      : N/A
BPDU Encap     : Dot1d
Active Protocol : N/A
Designated Port Id: 0

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

```

Number of SDPs : 1

* indicates that the corresponding row element may have been truncated.

A:Dut-A#

Table 60 Show Service-ID SDP Field Descriptions

Label	Description
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
Split Horizon Group	Name of the split horizon group that the SDP belongs to.
VC Type	Displays the VC type: ether or vlan.
VC Tag	Displays the explicit dot1Q value used when encapsulating to the SDP far end.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case).
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by 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.

Table 60 Show Service-ID SDP Field Descriptions (Continued)

Label	Description (Continued)
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. If the SDP type is GRE, then the following message displays: SDP delivery mechanism is not MPLS.

sdp

Syntax **sdp** *sdp-id* **pw-port** [*pw-port-id*]
sdp *sdp-id* **pw-port**
sdp *sdp-id* **pw-port** *pw-port-id* [**statistics**]
sdp [**consistent** | **inconsistent** | **na**] **egressifs**
sdp *sdp-id* **keep-alive-history**
sdp **far-end** {*ip-address* | *ipv6-address*} **keep-alive-history**
sdp [*sdp-id*] [**detail**]
sdp **far-end** {*ip-address* | *ipv6-address*} [**detail**]

Context show>service

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* — Specifies the SDP ID for which to display information.

Values 1 to 17407

Default All SDPs.

pw-port-id — Specifies the pseudo-wire port identifier.

Values 1 to 10239

statistics — Displays SDP statistics information.

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.

na — Indicates that there is no egress network interface that can carry traffic on this SDP.

egressifs — Indicates whether all the egress network interfaces that can carry traffic on this SDP are associated with the network-domain configured on this SDP.

ip-address | ipv6-address — Displays only SDPs matching with the specified far-end IP address. 64 characters maximum.

Default SDPs with any far-end IP address.

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

Default SDP summary output.

detail — Displays detailed SDP information.

Default SDP summary output.

Output The following output is an example of SDP information.

Sample Output

```
*A:ALA-12>config>service# show service sdp 1 pw-port
=====
Service Destination Point (sdp Id 1 Pw-Port)
=====
Pw-port   VC-Id   Adm    Encap   Opr    VC Type   Egr      Monitor
              Shaper  VPort   Oper
              Group
-----
1          1       up     dot1q   up     ether
2          2       up     qinq    up     ether
3          3       up     dot1q   up     ether
4          4       up     qinq    up     ether
-----
Entries found : 4
=====

*A:ALA-12>config>service# show service sdp 1 pw-port 3
=====
Service Destination Point (Sdp Id 1 Pw-Port 3)
=====
SDP Binding port      : lag-1
VC-Id                 : 3                      Admin Status          : up
Encap                 : dot1q                     Oper Status           : up
VC Type               : ether
Oper Flags            : (Not Specified)
Monitor Oper-Group    : (Not Specified)
=====

*A:ALA-12>config>service# show service sdp 1 pw-port 3 statistics
=====
Service Destination Point (Sdp Id 1 Pw-Port 3)
=====
SDP Binding port      : lag-1
VC-Id                 : 3                      Admin Status          : up
Encap                 : dot1q                     Oper Status           : up
VC Type               : ether
Oper Flags            : (Not Specified)
Monitor Oper-Group    : (Not Specified)
```

```

Statistics          :
I. Fwd. Pkts.       : 0
I. Fwd. Octs.       : 0
E. Fwd. Pkts.       : 0
I. Dro. Pkts.       : 0
I. Dro. Octs.       : 0
E. Fwd. Octets      : 0
=====

```

The following output is an example of SDP information, and [Table 61](#) describes the output fields.

Sample Output

```

*A:ALA-12# show service sdp
=====
Services: Service Destination Points
=====
SdpId    Adm MTU  Opr MTU  IP address    Adm  Opr      Deliver Signal
-----
10       4462    4462     10.20.1.3     Up   Dn NotReady MPLS   TLDP
40       4462    1534     10.20.1.20    Up   Up        MPLS   TLDP
60       4462    1514     10.20.1.21    Up   Up        GRE    TLDP
100      4462    4462     203.0.0.2     Down Down      GRE    TLDP
500      4462    4462     10.20.1.50    Up   Dn NotReady GRE    TLDP
-----
Number of SDPs : 5
=====
*A:ALA-12#

*A:ALA-12# show service sdp 2 detail
=====
Service Destination Point (Sdp Id : 2) Details
=====
-----
Sdp Id 2  -(10.10.10.104)
-----
Description          : GRE-10.10.10.104
SDP Id              : 2
Admin Path MTU       : 0
Oper Path MTU        : 0
Far End              : 10.10.10.104
Delivery              : GRE
Admin State          : Up
Oper State            : Down
Flags                : SignalingSessDown TransportTunnDown
Signaling             : TLDP
VLAN VC Etype        : 0x8100
Last Status Change   : 02/01/2007 09:11:39
Adv. MTU Over.       : No
Last Mgmt Change     : 02/01/2007 09:11:46

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

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
=====
*A:ALA-12#

```

```

*A:Dut-B# show service sdp
=====
Services: Service Destination Points
=====
SdpId  AdmMTU  OprMTU  Far End      Adm  Opr      Del      LSP  Sig
-----
230    0       1582    10.20.1.3    Up   Up        MPLS     I    TLDP
-----
Number of SDPs : 1
-----
Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable
=====
*A:Dut-B#
*A:Dut-B# show service sdp detail
=====
Services: Service Destination Points Details
=====
-----
Sdp Id 230 -10.20.1.3
-----
Description      : (Not Specified)
SDP Id           : 230
Admin Path MTU   : 0
Delivery         : MPLS
Far End          : 10.20.1.3
Tunnel Far End   : n/a
SDP Source       : manual
Oper Path MTU    : 1582
LSP Types        : SR-ISIS

Admin State      : Up
Signaling        : TLDP
Acct. Pol        : None
Last Status Change : 01/28/2015 22:00:07
Last Mgmt Change  : 01/28/2015 21:59:53
Bw BookingFactor : 100
Oper Max BW(Kbps) : 0
Net-Domain       : default
Flags            : None
Oper State       : Up
Metric           : 0
Collect Stats    : Disabled
Adv. MTU Over.   : No
VLAN VC Etype    : 0x8100
PBB Etype        : 0x88e7
Avail BW(Kbps)   : 0
Egr Interfaces   : Consistent

Mixed LSP Mode Information :
Mixed LSP Mode             : Disabled
Active LSP Type            : SR-ISIS

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

Src B-MAC LSB             : <none>
Ctrl PW Active            : n/a
Ctrl PW VC ID             : <none>

-----
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

```

```

-----
Segment Routing
-----

```

```

ISIS                  : enabled          LSP Id              : 524289
Oper Instance Id      : 0

```

```

Number of SDPs : 1

```

```

=====
*A:Dut-B#

```

```

*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#

```

Table 61 Show Service SDP Field Descriptions

Label	Description
SDP Id	The SDP identifier.
Adm MTU	Specifies the largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Opr MTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
IP address	Specifies the IP address of the remote end of the GRE or 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: GRE or MPLS.
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Hello Timeout	Specifies the number of seconds to wait for an SDP echo response message before declaring a timeout.
Unmatched Replies	Specifies the number of SDP unmatched message replies.
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.

Table 61 Show Service SDP Field Descriptions (Continued)

Label	Description (Continued)
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. If the SDP type is GRE, then the following message displays: SDP delivery mechanism is not MPLS.

subscriber-hosts

Syntax	subscriber-hosts [sap <i>sap-id</i>] [ip <i>ip-prefix/prefix-length</i>] [mac <i>ieee-address</i>] [sub-profile <i>sub-profile-name</i>] [sla-profile <i>sla-profile-name</i>] [app-profile <i>app-profile-name</i>] [wholesaler <i>service-id</i>] [address-origin <i>address-origin</i>] [detail]
Context	show>service>id
Description	This command displays subscriber host information.
Parameters	<i>sap-id</i> — Displays the specified subscriber host SAP information. <i>ip-prefix/prefix-length</i> — 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-le: 0 to 32 - ipv6-prefix: - x:x:x:x:x:x:x (eight 16-bit pieces) - x:x:x:x:x:x:d.d.d.d - x: [0 to FFFF] H - d: [0 to 255] D - ipv6-prefix-le: 0 to 128 <i>ieee-address</i> — 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. Values - xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx (cannot be all zeros) <i>sub-profile-name</i> — Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the config>subscr-mgmt>sub-profile context. 32 characters maximum. <i>sla-profile-name</i> — Specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the config>subscr-mgmt>sla-profile context. 32 characters maximum.

app-profile-name — Specifies an existing application profile name to be associated with the static subscriber host. The application profile is configured in the **config>subscr-mgmt>sla-profile** context. 32 characters maximum.

service-id — The VPRN service ID of the wholesaler (applies only to the 7750 SR).

Values 1 to 2147483648, *svc-name: 64 chars max*

address-origin — Shows the origin of the IP address assigned to the subscriber-host.

Values aaa, dynamic, static, bonding

detail — Displays detailed information.

log

Syntax **log**

Context show>service>id

Description This command displays event and accounting policy log information.

filter-id

Syntax **filter-id** [*filter-id*]

Context show>service>id>log

Description This command displays event file log information.

If no command line parameters are specified, a summary output of all event log files is displayed.

Specifying a file ID displays detailed information on the event file log.

Parameters *filter-id* — Specifies the filter policy.

Values 1 to 65535

Output [Table 62](#) describes the output fields for a log file summary.

Table 62 Filter-ID Field Descriptions

Label	Description
file-id	The log file ID.
rollover	The rollover time for the log file which is how long in between partitioning of the file into a new file.

Table 62 Filter-ID Field Descriptions (Continued)

Label	Description (Continued)
retention	The retention time for the file in the system which is how long the file should be retained in the file system.
admin location	The primary flash device specified for the file location. n/a — Indicates no specific flash device was specified.
backup location	The secondary flash device specified for the file location if the admin location is not available. n/a — Indicates that no backup flash device was specified.
oper location	The actual flash device on which the log file exists.
file-id	The log file ID.
rollover	The rollover time for the log file which is how long in between partitioning of the file into a new file.
retention	The retention time for the file in the system which is how long the file should be retained in the file system.
file name	The complete pathname of the file associated with the log ID.
expired	Indicates whether or not the retention period for this file has passed.
state	In progress — Indicates the current open log file. Complete — Indicates the old log file.

log-id

Syntax **log-id** [*log-id*] [**severity** *severity-level*] [**application** *application*] [**sequence** *from-seq* [*to-seq*]] [**count** *count*] [**subject** *subject* [**regexp**]] [**ascending** | **descending**] [**message** *format* [**msg-regexp**]]

Context show>service>id>log

Description This command displays an event log summary with settings and statistics or the contents of a specific log file, SNMP log, or memory log.

If the command is specified with no command line options, a summary of the defined system logs is displayed. The summary includes log settings and statistics.

If the log ID of a memory, SNMP, or file event log is specified, the command displays the contents of the log. Additional command line options control what and how the contents are displayed.

Contents of logs with console, session or syslog destinations cannot be displayed. The actual events can only be viewed on the receiving syslog or console device.

Parameters *log-id* — Displays the contents of the specified file log or memory log ID. The log ID must have a destination of an SNMP or file log or a memory log for this parameter to be used.

Values 1 to 100

Default Displays the event log summary.

severity-level — Displays only events with the specified and higher severity.

Values cleared, indeterminate, critical, major, minor, warning

Default All severity levels.

application — Displays only events generated by the specified application.

The following values apply to the 7750 SR and 7950 XRS only:

Values application_assurance, aps, atm, bfd, bgp, calltrace, ccag, cflowd, chassis, cpmhwfilter, cpmhwqueue, debug, dhcp, dhcps, diameter, dot1x, dynsvc, efm_oam, elmi, ering, eth_cfm, etun, filter, fpe, gsmp, gmpls, gtungrp, igh, igmp, igmp_snooping, ip, ipfix, ipsec, ipsec_cpm, isis, l2tp, lag, ldap, ldp, li, lldp, lmp, logger, mcac, mcpaath, mc_redundancy, mirror, mld, mld_snooping, mpls, mpls_tp, mrp, msdp, nat, ntp, oam, open_flow, ospf, pcep, pim, pim_snooping, port, ppp, pppoe, pppoe_clnt, ptp, pxc, python, qos, radius, rip, rip_ng, route_next_hop, route_policy, rpki, rsvp, security, sflow, snmp, stp, subscr_mgmt, sub_host_trk, svcmgr, system, tip, tls, user, user_db, video, vrrp, vrtr, wlan_gw, wpp

The following values apply only to the 7450 ESS:

Values application_assurance, aps, atm, bfd, bgp, calltrace, cflowd, chassis, dhcp, debug, dhcps, diameter, dynsvc, efm_oam, elmi, ering, eth_cfm, etun, filter, gsmp, gmpls, igh, igmp, igmp_snooping, ip, ipsec, isis, l2tp, lag, ldap, ldp, li, lldp, logger, mcpaath, mc_redundancy, mirror, mld, mld_snooping, mpls, mpls_tp, msdp, nat, oam, open_flow, ospf, pim, pim_snooping, port, ppp, pppoe, ptp, radius, rip, rip_ng, route_policy, rpki, rsvp, security, snmp, stp, svcmgr, system, user, video, vrrp, vrtr, wlan_gw, wpp

Default All applications.

from-seq [to-seq] — Displays the log entry numbers from a particular entry sequence number (*from-seq*) to another sequence number (*to-seq*). The *to-seq* value must be larger than the *from-seq* value.

If the *to-seq* number is not provided, the log contents to the end of the log is displayed unless the **count** parameter is present in which case the number of entries displayed is limited by the **count**.

Values 1 to 4294967295

Default All sequence numbers.

count — Limits the number of log entries displayed to the *number* specified.

Values 1 to 4294967295

Default All log entries.

subject — Displays only log entries matching the specified text *subject* string. The subject is the object affected by the event, for example the port-id would be the subject for a link-up or link-down event. 32 characters maximum.

regexp — Specifies to use a regular expression as parameters with the specified **subject** string.

ascending | descending — Specifies sort direction. Logs are normally shown from the newest entry to the oldest in **descending** sequence number order on the screen. When using the **ascending** parameter, the log will be shown from the oldest to the newest entry.

Default Descending.

format — Specifies a message string up to 400 characters to be used in the display criteria.

msg-regexp — Specifies to use a regular expression as parameters with the specified message string.

Output [Table 63](#) describes the log ID field output.

Table 63 Show Log-ID Field Descriptions

Label	Description
Log Id	An event log destination.
Source	no — The event log filter is not currently in use by a log ID. yes — The event log filter is currently in use by a log ID.
Filter ID	The value is the index to the entry which defines the filter to be applied to this log's source event stream to limit the events output to this log's destination. If the value is 0, then all events in the source log are forwarded to the destination.
Admin State	Up — Indicates that the administrative state is up. Down — Indicates that the administrative state is down.
Oper State	Up — Indicates that the operational state is up. Down — Indicates that the operational state is down.
Logged	The number of events that have been sent to the log source(s) that were forwarded to the log destination.
Dropped	The number of events that have been sent to the log source(s) that were not forwarded to the log destination because they were filtered out by the log filter.

Table 63 Show Log-ID Field Descriptions (Continued)

Label	Description (Continued)
Dest. Type	Console — All selected log events are directed to the system console. If the console is not connected, then all entries are dropped. Syslog — All selected log events are sent to the syslog address. SNMP traps — Events defined as SNMP traps are sent to the configured SNMP trap destinations and are logged in NOTIFICATION-LOG-MIB tables. File — All selected log events will be directed to a file on one of the CPM's compact flash disks. (7750 SR and 7450 ESS only). Memory — All selected log events will be directed to an in-memory storage area.
Dest ID	The event log stream destination.
Size	The allocated memory size for the log.
Time format	The time format specifies the type of timestamp format for events sent to logs where log ID destination is either syslog or file. When the time format is UTC, timestamps are written using the Coordinated Universal Time value. When the time format is local, timestamps are written in the system's local time.

snmp-trap-group

Syntax `snmp-trap-group [log-id]`

Context `show>service>id>log`

Description This command displays SNMP trap group configuration information.

Parameters *log-id* — Displays only SNMP trap group information for the specified trap group log ID.

Values 1 to 100

Output **SNMP Trap Group Output**

[Table 64](#) describes SNMP trap group output fields.

Table 64 SNMP Trap Group Field Descriptions

Label	Description
Log-ID	The log destination ID for an event stream.

Table 64 SNMP Trap Group Field Descriptions (Continued)

Label	Description
Address	The IP address of the trap receiver.
Port	The destination UDP port used for sending traps to the destination, expressed as a decimal integer.
Version	Specifies the SNMP version format to use for traps sent to the trap receiver. Valid values are snmpv1, snmpv2c, snmpv3.
Community	The community string required by snmpv1 or snmpv2c trap receivers.
Security-Level	The required authentication and privacy levels required to access the views on this node.
Replay	Indicates whether or not the replay parameter has been configured, enabled or disabled, for the trap-target address.
Replay from	Indicates the sequence ID of the first missed notification that will be replayed when a route is added to the routing table by which trap-target address can be reached. If no notifications are waiting to be replayed this field shows n/a.
Last Replay	Indicates the last time missed events were replayed to the trap-target address. If no events have ever been replayed this field shows never.

syslog

Syntax **syslog** [*syslog-id*]

Context show>service>id>log

Description This command displays syslog event log destination summary information or detailed information on a specific syslog destination.

Parameters *syslog-id* — Displays detailed information on the specified syslog event log destination.

Values 1 to 10

Output **Syslog Event Log Destination Summary Output**

[Table 65](#) describes the syslog output fields.

Table 65 Show Log Syslog Field Descriptions

Label	Description
Syslog ID	The syslog ID number for the syslog destination.
IP Address	The IP address of the syslog target host.
Port	The configured UDP port number used when sending syslog messages.
Facility	The facility code for messages sent to the syslog target host.
Severity Level	The syslog message severity level threshold.
Below Level Dropped	A count of messages not sent to the syslog collector target because the severity level of the message was above the configured severity. The higher the level, the lower the severity.
Prefix Present	Yes — A log prefix was prepended to the syslog message sent to the syslog host. No — A log prefix was not prepended to the syslog message sent to the syslog host.
Description	A text description stored in the configuration file for a configuration context.
LogPrefix	The prefix string prepended to the syslog message.
Log-id	Events are directed to this <i>destination</i> .

aggregate

Syntax **aggregate** [*family*] [**active**] [**detail**]

Context show>router

Description This command displays aggregated routes.

Parameters **active** — This keyword filters out inactive aggregates.

detail — This keyword displays detailed information.

family — Specifies the family to display.

Values ipv4, ipv6

Output The following output is an example of router aggregate information, and [Table 66](#) 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#

*A:Dut-A>config>router# show router aggregate

=====
Aggregates (Router: Base)
=====
Prefix                Summary  AS Set   Aggr IP-Address  Aggr AS
                    NextHop                               State
                    NextHopType
-----
10.2.3.0/24           False   0.0.0.0   0
                    10.2.2.2           Inactive
                    Indirect
10.2.0.0/16           False   0.0.0.0   0
                    False           Active
                    None
-----
No. of Aggregates: 2
=====

*A:CPM133>config>router# show router aggregate

=====
Aggregates (Router: Base)
=====
Prefix                Summary  AS Set   Aggr IP-Address  Aggr AS
                    NextHop                               State
                    NextHopType
-----
10.0.0.0/8           False   0.0.0.0   0
                    False           Inactive
                    100:33          Blackhole
-----
No. of Aggregates: 1
=====

```

Table 66 Show Aggregate Field Descriptions

Label	Description
Prefix	Displays the destination address of the aggregate route in dotted decimal notation.

Table 66 Show Aggregate Field Descriptions (Continued)

Label	Description
Summary	Specifies whether the aggregate or more specific components are advertised.
AS Set	Displays an aggregate where the path advertised for the route consists of all elements contained in all paths that are being summarized.
Aggr AS	Displays the aggregator path attribute to the aggregate route.
Aggr IP-Address	The IP address of the aggregated route.
State	The operational state of the aggregated route.
No. of Aggregates	The total number of aggregated routes.

arp

Syntax **arp** [*ip-int-name* | **ip-address** [/mask] | **mac** *ieee-mac-addr* | **summary**] [*arp-type*]

Context show>router

Description This command displays the router ARP table sorted by IP address.

If no command line options are specified, all ARP entries are displayed.

Parameters *ip-address[/mask]* — Only displays ARP entries associated with the specified IP address.

Values ip-address: a.b.c.d
mask: 1 to 32

ip-int-name — Only displays ARP entries associated with the specified IP interface name. 32 characters maximum.

ieee-mac-addr — Only displays ARP entries associated with the specified MAC address.

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

summary — Displays an abbreviated list of ARP entries.

arp-type — Only displays ARP entries of the specific type.

Values local, dynamic, static, managed

Output The following output is an example of ARP table information, and [Table 67](#) 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 Oth      system
10.10.4.3       00:00:00:00:00:00 00h00m00s Oth      ALA-1-2
10.10.5.3       00:00:00:00:00:00 00h00m00s Oth      ALA-1-3
10.10.7.3       00:00:00:00:00:00 00h00m00s Oth      ALA-1-5
10.10.0.16      00:00:00:00:00:00 00h00m00s Oth      bozo
10.10.3.3       00:00:00:00:00:00 00h00m00s Oth      gizmo
10.10.2.3       00:00:00:00:00:00 00h00m00s Oth      hobo
10.10.1.17      00:00:00:00:00:00 00h00m00s Oth      int-cflowd
10.0.0.92       00:00:00:00:00:00 04h00m00s Dyn      to-104
10.0.0.103      04:67:01:01:00:01 00h00m00s Oth[I]   to-104
10.0.0.104      04:68:01:01:00:01 03h59m49s Dyn[I]   to-104
10.10.36.2      00:00:00:00:00:00 00h00m00s Oth      tuesday
192.168.2.98    00:03:47:c8:b4:86 00h14m37s Dyn[I]   management
192.168.2.103   00:03:47:dc:98:1d 00h00m00s Oth[I]   management
-----
No. of ARP Entries: 14
=====
*A:ALA-12#

*A:ALA-12# show router 3 arp 10.10.0.3
=====
ARP Table
=====
IP Address      MAC Address      Expiry      Type      Interface
-----
10.10.0.3       04:5d:ff:00:00:00 00:00:00 Oth      system
=====
*A:ALA-12#

*A:ALA-12# show router 3 arp to-ser1
=====
ARP Table
=====
IP Address      MAC Address      Expiry      Type      Interface
-----
10.10.13.1      04:5b:01:01:00:02 03:53:09 Dyn      to-ser1
=====
*A:ALA-12#

```

Table 67 ARP Table Field Descriptions

Label	Description
IP Address	The IP address of the ARP entry.
MAC Address	The MAC address of the ARP entry.
Expiry	The age of the ARP entry.

Table 67 ARP Table Field Descriptions (Continued)

Label	Description (Continued)
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* [/i>ip-prefix-length]] [**damp-type**] [**detail**] [**ipv4**]
damping [*ip-prefix* [/i>ip-prefix-length]] [**damp-type**] [**detail**] {**ipv6** | **label-ipv4** | **label-ipv6** |
mcast-ipv4 | **mcast-ipv6** | **mvpn-ipv4** | **vpn-ipv4** | **vpn-ipv6**}

Context show>router>bgp

Description This command displays BGP routes with have been dampened due to route flapping. This command can be entered with or without a route parameter.

When the keyword **detail** is included, more detailed information displays.

When only the command is entered (without any parameters included except **detail**), then all dampened routes are listed.

When a parameter is specified, then the matching route or routes are listed.

When a **decayed**, **history**, or **suppressed** keyword is specified, only those types of dampened routes are listed.

Parameters *ip-prefix*[*ip-prefix-length*] — Displays damping information for the specified IP prefix length.

Values ipv4-prefix: a.b.c.d (host bits must be 0)
 ipv4-prefix-le: 0 to 32
 ipv6-prefix:
 • x:x:x:x:x:x:x (eight 16-bit pieces)
 • x:x:x:x:x:d.d.d.d
 • x: [0 to FFFF] H
 • d: [0 to 255] D
 ipv6-prefix-le: 0 to 128

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.

ipv4 — Displays damped routes for the IPv4 family.

ipv6 — Displays damped routes for the IPv6 family.

label-ipv4 — Displays damped routes for the label IPv4 family.

label-ipv6 — Displays damped routes for the label IPv6 family.

mcast-ipv4 — Displays damped routes for the MCAST IPv4 family.

mcast-ipv6 — Displays damped routes for the MCAST IPv6 family.

mvpn-ipv4 — Displays damped routes for the MVPN IPv4 family.

vpn-ipv4 — Displays damped routes for the VPN IPv4 family.

vpn-ipv6 — Displays damped routes for the VPN IPv6 family.

Output The following output is an example of BGP damping information, and [Table 68](#) describes the output fields.

Sample Output

```
*A:ALA-12# show router 3 bgp damping
=====
BGP Router ID : 10.0.0.14          AS : 65206    Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes
=====
Flag  Network          From          Reuse          AS-Path
-----
ud*i  10.149.7.0/24       10.0.28.1     00h00m00s      60203 65001 19855 3356
                        1239  22406
si    10.155.6.0/23      10.0.28.1     00h43m41s      60203 65001 19855 3356
                        2914  7459
si    10.155.8.0/22      10.0.28.1     00h38m31s      60203 65001 19855 3356
                        2914  7459
si    10.155.12.0/22     10.0.28.1     00h35m41s      60203 65001 19855 3356
                        2914  7459
si    10.155.22.0/23     10.0.28.1     00h35m41s      60203 65001 19855 3356
                        2914  7459
si    10.155.24.0/22     10.0.28.1     00h35m41s      60203 65001 19855 3356
                        2914  7459
si    10.155.28.0/22     10.0.28.1     00h34m31s      60203 65001 19855 3356
                        2914  7459
si    10.155.40.0/21     10.0.28.1     00h28m24s      60203 65001 19855 3356
                        7911  7459
si    10.155.48.0/20     10.0.28.1     00h28m24s      60203 65001 19855 3356
                        7911  7459
ud*i  10.8.140.0/24       10.0.28.1     00h00m00s      60203 65001 19855 3356
```

```

                                4637 17447
ud*i 10.8.141.0/24      10.0.28.1      00h00m00s 60203 65001 19855 3356
                                4637 17447
ud*i 10.9.0.0/18      10.0.28.1      00h00m00s 60203 65001 19855 3356
                                3561 9658 6163
. . .
ud*i 10.213.184.0/23  10.0.28.1      00h00m00s 60203 65001 19855 3356
                                6774 6774 9154
-----
*A:ALA-12#

*A:ALA-12# show router 3 bgp damping detail
=====
BGP Router ID : 10.0.0.14      AS : 65206      Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * -
valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes
=====
-----
Network : 10.149.7.0/24
-----
Network      : 10.149.7.0/24      Peer      : 10.0.28.1
NextHop      : 10.0.28.1      Reuse time : 00h00m00s
Peer AS      : 60203      Peer Router-Id : 10.32.27.203
Local Pref   : none
Age          : 00h22m09s      Last update : 02d00h58m
FOM Present  : 738      FOM Last upd. : 2039
Number of Flaps : 2      Flags      : ud*i
Path         : 60203 65001 19855 3356 1239 22406
Applied Policy : default-damping-profile
-----
Network : 10.142.48.0/20
-----
Network      : 10.142.48.0/20      Peer      : 10.0.28.1
NextHop      : 10.0.28.1      Reuse time : 00h00m00s
Peer AS      : 60203      Peer Router-Id : 10.32.27.203
Local Pref   : none
Age          : 00h00m38s      Last update : 02d01h20m
FOM Present  : 2011      FOM Last upd. : 2023
Number of Flaps : 2      Flags      : ud*i
Path         : 60203 65001 19855 3356 3561 5551 1889
Applied Policy : default-damping-profile
-----
Network : 10.200.128.0/19
-----
Network      : 10.200.128.0/19      Peer      : 10.0.28.1
NextHop      : 10.0.28.1      Reuse time : 00h00m00s
Peer AS      : 60203      Peer Router-Id : 10.32.27.203
Local Pref   : none
Age          : 00h00m38s      Last update : 02d01h20m
FOM Present  : 2011      FOM Last upd. : 2023
Number of Flaps : 2      Flags      : ud*i
Path         : 60203 65001 19855 1299 702 1889
Applied Policy : default-damping-profile

```

```
-----
Network : 10.203.192.0/18
-----
Network      : 10.203.192.0/18      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h00m00s
Peer AS      : 60203              Peer Router-Id : 10.32.27.203
Local Pref   : none
Age          : 00h00m07s          Last update : 02d01h20m
FOM Present  : 1018              FOM Last upd. : 1024
Number of Flaps : 1              Flags       : ud*i
Path         : 60203 65001 19855 1299 702 1889
Applied Policy : default-damping-profile
-----
*A:ALA-12#

*A:ALA-12# show router 3 bgp damping 10.203.192.0/18 detail
=====
BGP Router ID : 10.0.0.14      AS : 65206      Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes 10.203.192.0/18
=====
Network : 10.203.192.0/18
-----
Network      : 10.203.192.0/18      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h00m00s
Peer AS      : 60203              Peer Router-Id : 10.32.27.203
Local Pref   : none
Age          : 00h00m42s          Last update : 02d01h20m
FOM Present  : 2003              FOM Last upd. : 2025
Number of Flaps : 2              Flags       : ud*i
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Paths : 1
=====
*A:ALA-12#
*A:ALA-12# show router 3 bgp damping suppressed detail
=====
BGP Router ID : 10.0.0.14      AS : 65206      Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes (Suppressed)
=====
Network : 10.142.48.0/20
-----
Network      : 10.142.48.0/20      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h29m22s
Peer AS      : 60203              Peer Router-Id : 10.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update : 02d01h20m
FOM Present  : 2936              FOM Last upd. : 3001
```

```

Number of Flaps : 3                      Flags : si
Path           : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Network : 10.200.128.0/19
-----
Network      : 10.200.128.0/19      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h29m22s
Peer AS      : 60203              Peer Router-Id : 10.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update : 02d01h20m
FOM Present  : 2936              FOM Last upd. : 3001
Number of Flaps : 3              Flags      : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Network : 10.203.240.0/20
-----
Network      : 10.203.240.0/20      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h29m22s
Peer AS      : 60203              Peer Router-Id : 10.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update : 02d01h20m
FOM Present  : 2936              FOM Last upd. : 3001
Number of Flaps : 3              Flags      : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Network : 10.206.0.0/17
-----
Network      : 10.206.0.0/17      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h29m22s
Peer AS      : 60203              Peer Router-Id : 10.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update : 02d01h20m
FOM Present  : 2936              FOM Last upd. : 3001
Number of Flaps : 3              Flags      : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
*A:ALA-12#

```

Table 68 Show Damping Field Descriptions

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured or inherited local AS for the specified peer group. If not configured, then it is the same value as the AS.
Network	Route IP prefix and mask length for the route.

Table 68 Show Damping Field Descriptions (Continued)

Label	Description (Continued)
Flag(s)	Legend: Status codes: u- used, s-suppressed, h-history, d-decayed, *-valid. If a * is not present, then the status is invalid. Origin codes: i-IGP, e-EGP, ?-incomplete, >-best
Network	The IP prefix and mask length for the route.
From	The originator ID path attribute value.
Reuse time	The time when a suppressed route can be used again.
AS Path	The BGP AS path for the route.
Peer	The router ID of the advertising router.
NextHop	BGP nexthop for the route.
Peer AS	The autonomous system number of the advertising router.
Peer Router-Id	The router ID of the advertising router.
Local Pref	BGP local preference path attribute for the route.
Age	The time elapsed since the service was enabled.
Last update	The time when BGP was updated last in second/minute/hour (SS:MM:HH) format.
FOM Present	The current Figure of Merit (FOM) value.
Number of Flaps	The number of flaps in the neighbor connection.
Reuse time	The time when the route can be reused.
Path	The BGP AS path for the route.
Applied Policy	The applied route policy name.

group

Syntax **group** [*name*] [*detail*]

Context show>router>bgp

Description This command displays group information for a BGP peer group. This command can be entered with or without parameters.

When this command is entered without a group name, information about all peer groups displays.

When the command is issued with a specific group name, information only pertaining to that specific peer group displays.

The 'State' field displays the BGP group's operational state. Other valid states are:

- Up - BGP global process is configured and running.
- Down - BGP global process is administratively shutdown and not running.
- Disabled - BGP global process is operationally disabled. The process must be restarted by the operator.

Parameters *name* — Displays information for the BGP group specified.

detail — Displays detailed information.

Output The following output is an example of BGP group information, and [Table 69](#) describes the d output fields.

Sample Output (7450 ESS and 7750 SR)

```
*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
NLRI             : Unicast
Keep Alive       : 30
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 69 Standard and Detailed Group Field Descriptions

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.

Table 69 **Standard and Detailed Group Field Descriptions (Continued)**

Label	Description (Continued)
State	Disabled — The BGP peer group has been operationally disabled. Down — The BGP peer group is operationally inactive. Up — The BGP peer group is operationally active.
Peer AS	The configured or inherited peer AS for the specified peer group.
Local AS	The configured or inherited local AS for the specified peer group.
Local Address	The configured or inherited local address for originating peering for the specified peer group.
Loop Detect	The configured or inherited loop detect setting for the specified peer group.
Connect Retry	The configured or inherited connect retry timer value.
Authentication	n/a — No authentication is configured. MD5 — MD5 authentication is configured.
Local Pref	The configured or inherited local preference value.
MED Out	The configured or inherited MED value assigned to advertised routes without a MED attribute.
Min Route Advt.	The minimum amount of time that must pass between route updates for the same IP prefix.
Min AS Originate	The minimum amount of time that must pass between updates for a route originated by the local router.
Multihop	The maximum number of router hops a BGP connection can traverse.
Multipath	The configured or inherited multipath value, determining the maximum number of ECMP routes BGP can advertise to the RTM.
Prefix Limit	No Limit — No route limit assigned to the BGP peer group. 1 to 4294967295 — The maximum number of routes BGP can learn from a peer.
Passive	Disabled — BGP attempts to establish BGP connections with neighbors in the specified peer group. Enabled — BGP will not actively attempt to establish BGP connections with neighbors in the specified peer group.
Next Hop Self	Disabled — BGP is not configured to send only its own IP address as the BGP nexthop in route updates to neighbors in the peer group. Enabled — BGP sends only its own IP address as the BGP nexthop in route updates to neighbors in the specified peer group.

Table 69 Standard and Detailed Group Field Descriptions (Continued)

Label	Description (Continued)
Aggregator ID 0	Disabled — BGP is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group. Enabled — BGP is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group.
Remove Private	Disabled — BGP will not remove all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group. Enabled — BGP removes all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group.
Damping	Disabled — The peer group is configured not to dampen route flaps. Enabled — The peer group is configured to dampen route flaps.
Export Policy	The configured export policies for the peer group.
Import Policy	The configured import policies for the peer group.
Hold Time	The configured hold time setting.
Keep Alive	The configured keepalive setting.
Cluster Id	n/a — 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 [family] filter2*
neighbor *[as-number [family] filter2*

Context show>router>bgp

Description This command displays BGP neighbor information. This command can be entered with or without any parameters.

When this command is issued without any parameters, information about all BGP peers displays.

When the command is issued with a specific IP address or ASN, information regarding only that specific peer or peers with the same AS display.

When either **received-routes** or **advertised-routes** is specified, then the routes received from or sent to the specified peer is listed (see second output example). This information is not available by SNMP.

When either **history** or **suppressed** is specified, then the routes learned from those peers that either have a history or are suppressed (respectively) are listed.

The 'State' field displays the BGP peer's protocol state. In addition to the standard protocol states, this field can also display the 'Disabled' operational state which indicates the peer is operationally disabled and must be restarted by the operator.

- Parameters**
- ip-addr* — Displays the BGP neighbor with the specified IP address.
 - family* — Specifies the type of routing information to be distributed by the BGP instance.
 - Values** ipv4, vpn-ipv4, ipv6, mcast-ipv4, vpn-ipv6, l2-vpn, mvpn-ipv4, mdt-safi, flow-ipv4, ms-pw, route-target, mcast-vpn-ipv4, mvpn-ipv6, flow-ipv6, evpn, mcast-ipv6, label-ipv4, label-ipv6
 - filter2* — Specifies route criteria.
 - Values** history, suppressed

Output The following output is an example of BGP neighbor information, and [Table 70](#) describes the output fields.

Sample Output

```
*A:ALA-12# show router 3 bgp neighbor
=====
BGP Neighbor
=====
-----
Peer : 10.0.0.15          Group : To_AS_40000
-----
-----
Peer AS      : 65205
Peer Address : 10.0.0.15      Peer Port    : 0
Local AS     : 65206
Local Address : 10.0.0.16    Local Port    : 0
Peer Type    : External
State        : Active        Last State    : Connect
Last Event   : openFail
Last Error   : Hold Timer Expire
Hold Time    : 90
Active Hold Time : 0          Keep Alive    : 30
Cluster Id   : None          Active Keep Alive: 0
Preference   : 170
Recd. Prefixes : 0          Num of Flaps   : 0
Recd. Paths   : 0          Active Prefixes : 0
Input Queue   : 0          Suppressed Paths : 0
Output 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 70 **Standard and Detailed Neighbor Field Descriptions**

Label	Description
Peer	The IP address of the configured BGP peer.
Group	The BGP peer group to which this peer is assigned.
Peer AS	The configured or inherited peer AS for the peer group.
Peer Address	The configured address for the BGP peer.

Table 70 **Standard and Detailed Neighbor Field Descriptions (Continued)**

Label	Description (Continued)
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 displayed in addition, 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 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. Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION.

Table 70 Standard and Detailed Neighbor Field Descriptions (Continued)

Label	Description (Continued)
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. open — BGP transport connection opened. close — BGP transport connection closed. openFail — BGP transport connection failed to open. error — BGP transport connection error. connectRetry — Connect retry timer expired. holdTime — Hold time timer expired. keepAlive — Keepalive timer expired. recvOpen — Receive an OPEN message. revKeepalive — Receive an KEEPALIVE message. recvUpdate — Receive an UPDATE message. recvNotify — Receive an NOTIFICATION message. None — No events have occurred.
Last Error	Displays the last BGP error and sub-code to occur on the BGP neighbor.
Connect Retry	The configured or inherited connect retry timer value.
Local Pref.	The configured or inherited local preference value.
Min Route Advt.	The minimum amount of time that must pass between route updates for the same IP prefix.
Min AS Originate	The minimum amount of time that must pass between updates for a route originated by the local router.
Multihop	The maximum number of router hops a BGP connection can traverse.
Multipath	The configured or inherited multipath value, determining the maximum number of ECMP routes BGP can advertise to the RTM.
Damping	Disabled — BGP neighbor is configured not to dampen route flaps. Enabled — BGP neighbor is configured to dampen route flaps.
Loop Detect	Ignore — The BGP neighbor is configured to ignore routes with an AS loop. Drop — The BGP neighbor is configured to drop the BGP peering if an AS loop is detected. Off — AS loop detection is disabled for the neighbor.
MED Out	The configured or inherited MED value assigned to advertised routes without a MED attribute.

Table 70 **Standard and Detailed Neighbor Field Descriptions (Continued)**

Label	Description (Continued)
Authentication	None — No authentication is configured. MD5 — MD5 authentication is configured.
Next Hop Self	Disabled — BGP is not configured to send only its own IP address as the BGP nexthop in route updates to the specified neighbor. Enabled — BGP will send only its own IP address as the BGP nexthop in route updates to the neighbor.
AggregatorID Zero	Disabled — The BGP Neighbor is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates. Enabled — The BGP Neighbor is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates.
Remove Private	Disabled — BGP will not remove all private AS numbers from the AS path attribute, in updates sent to the specified neighbor. Enabled — BGP will remove all private AS numbers from the AS path attribute, in updates sent to the specified neighbor.
Passive	Disabled — BGP will actively attempt to establish a BGP connection with the specified neighbor. Enabled — BGP will not actively attempt to establish a BGP connection with the specified neighbor.
Prefix Limit	No Limit — No route limit assigned to the BGP peer group. 1 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. 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.

Table 70 Standard and Detailed Neighbor Field Descriptions (Continued)

Label	Description (Continued)
Active Prefixes	The number of routes received from the BGP neighbor and active in the forwarding table.
Recd. Paths	The number of unique sets of path attributes received from the BGP neighbor.
Suppressed Paths	The number of unique sets of path attributes received from the BGP neighbor and suppressed due to route damping.
Input Queue	The number of BGP messages to be processed.
Output Queue	The number of BGP messages to be transmitted.
i/p Messages	Total number of packets received from the BGP neighbor.
o/p Messages	Total number of packets sent to the BGP neighbor.
i/p Octets	Total number of octets received from the BGP neighbor.
o/p Octets	Total number of octets sent to the BGP neighbor.
i/p Updates	Total number of BGP updates received from the BGP neighbor.
o/p Updates	Total number of BGP updates sent to the BGP neighbor.
Export Policy	The configured export policies for the peer group.
Import Policy	The configured import policies for the peer group.

Show Advertised and Received Routes Output

The following output is an example of standard and detailed information for a neighbor, and [Table 71](#) describes the output fields.

Sample Output

```
*A:ALA-12# show router 3 bgp neighbor 10.0.0.16 received-routes
=====
BGP Router ID : 10.0.0.16      AS : 65206   Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP Neighbor
=====
Flag  Network          Nexthop          LocalPref  MED      As-Path
-----
?    10.0.0.16/32        10.0.0.16        100        none     No As-Path
?    10.0.6.0/24         10.0.0.16        100        none     No As-Path
?    10.0.8.0/24         10.0.0.16        100        none     No As-Path
?    10.0.12.0/24        10.0.0.16        100        none     No As-Path
```

```
?      10.0.13.0/24      10.0.0.16      100      none      No As-Path
?      10.0.204.0/24     10.0.0.16      100      none      No As-Path
=====
*A:ALA-12#
```

Table 71 Show Advertised and Received Routes Field Descriptions

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured local AS setting. If not configured, then 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.
Next Hop	BGP nexthop for the route.
LocalPref	BGP local preference path attribute for the route.
MED	BGP Multi-Exit Discriminator (MED) path attribute for the route.
AS Path	The BGP AS path for the route.

paths

- Syntax** `paths`
- Context** `show>router>bgp`
- Description** This command displays a summary of BGP path attributes.
- Output** the following output is an example of BGP path information, and [Table 72](#) 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 72 **Show Path Field Descriptions**

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured local AS setting. If not configured, then the value is the same as the AS.
Path	The AS path attribute.
Origin	EGP — The NLRI is learned by an EGP protocol. IGP — The NLRI is interior to the originating AS. INCOMPLETE — NLRI was learned another way.
Next Hop	The advertised BGP nexthop.
MED	The Multi-Exit Discriminator value.
Local Preference	The local preference value.
Refs	The number of routes using a specified set of path attributes.
ASes	The number of autonomous system numbers in the AS path attribute.
Segments	The number of segments in the AS path attribute.

Table 72 Show Path Field Descriptions (Continued)

Label	Description (Continued)
Flags	EBGP-learned — Path attributes learned by an EBGP peering. IBGP-Learned — Path attributes learned by an IBGP peering.
Aggregator	The route aggregator ID.
Community	The BGP community attribute list.
Originator ID	The originator ID path attribute value.
Cluster List	The route reflector cluster list.

routes

- Syntax** `routes [ip-prefix/mask | ip-address]`
- Context** `show>router>bgp`
- Description** This command displays BGP route information.
- When this command is issued without any parameters, then the entire BGP routing table displays.
- When this command is issued with an IP prefix/mask or IP address, then the best match for the parameter displays.
- Parameters** `ip-prefix/mask | ip-address` — Displays routes information for the specified IP prefix and mask, or IP address. 256 characters max
- Output** The following output is an example of BGP routes, and [Table 73](#) 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 192.0.0.0/31 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      : 192.0.0.0/31
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#

*A:Dut-B>config>service>vprn>bgp# show router bgp routes 10.5.5.5/32 hunt
=====
BGP Router ID:10.20.1.2      AS:1      Local AS:1
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
              l - leaked
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
RIB In Entries
-----
Network      : 10.5.5.5/32
Nexthop      : 10.1.1.2 (VPRN 1)
Path Id      : None
From         : ::
Res. Nexthop : 10.1.1.2
Local Pref.  : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric  : None
Connector    : None
Community    : 10:2
Interface Name : to_Dut-A
Aggregator    : None
MED           : None
```

```

Cluster      : No Cluster Members
Originator Id : None                      Peer Router Id : 0.0.0.0
Fwd Class    : None                      Priority       : None
Flags        : Used Valid Best IGP Leaked
Route Source  : Leaked from VPRN 1
AS-Path       : 2011
Route Tag     : 0
Neighbor-AS   : 2011
Orig Validation: NotFound
Source Class  : 0                        Dest Class     : 0
...

```

Table 73 Show BGP Routes Field Descriptions

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured local AS setting, if not configured it is the same as the system AS.
Network	The IP prefix and mask length.
Nexthop	The BGP nexthop.
From	The advertising BGP neighbor's IP address.
Res. Nexthop	The resolved nexthop.
Local Pref.	The local preference value.
Flag	u – used s – suppressed h – history d – decayed * – valid i – igp ? – 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.

Table 73 Show BGP Routes Field Descriptions (Continued)

Label	Description
MED	The MED metric value. none — No MED metric is present.
Community	The BGP community attribute list.
Cluster	The route reflector cluster list.
Originator Id	The originator ID path attribute value. none — The originator ID attribute is not present.
Peer Router Id	The router ID of the advertising router.
AS-Path	The BGP AS path attribute.
VPRN Imported	Displays the VPRNs where a particular BGP-VPN received route has been imported and installed.

summary

Syntax **summary** [**all**]
summary [**family** *family*] [**group** *name*]

Context show>router>bgp

Description This command displays a summary of BGP neighbor information.

If confederations are not configured, that portion of the output will not display.

The “State” field displays the global BGP operational state. The valid values are:

- Up — BGP global process is configured and running.
Down — BGP global process is administratively shutdown and not running.
Disabled — BGP global process is operationally disabled. The process must be restarted by the operator.
- For example, if a BGP peer is operationally disabled, then the state in the summary table shows the state ‘Disabled’

Parameters **all** — Displays BGP peers in all instances.

family — Specifies the type of routing information to be distributed by the BGP instance.

Values pv4, vpn-ipv4, ipv6, mcast-ipv4, vpn-ipv6, l2-vpn, mdt-safi, ms-pw, mvpn-ipv4, flow-ipv4, route-target, mcast-vpn-ipv4, mvpn-ipv6, flow-ipv6, evpn, mcast-ipv6, label-ipv4, label-ipv6, bgp-ls

name — Specifies the group name. 32 characters maximum

Output The following output is an example of BGP summary information, and [Table 74](#) describes the output fields.

Sample Output

```
*A:ALA-12# show router 3 bgp summary
=====
BGP Router ID : 10.0.0.14          AS : 65206    Local AS : 65206
=====
BGP Admin State      : Up          BGP Oper State      : Up
Confederation AS     : 40000
Member Confederations : 65205 65206 65207 65208

Number of Peer Groups : 2          Number of Peers      : 7
Total BGP Active Routes : 86689    Total BGP Routes     : 116999
Total BGP Paths        : 35860     Total Path Memory    : 2749476
Total Suppressed Routes : 0        Total History Routes  : 0
Total Decayed Routes   : 0

=====
BGP Summary
=====
Neighbor           AS PktRcvd PktSent InQ OutQ   Up/Down State|Recv/Actv/Sent
-----
10.0.0.1           65206      5  21849   0   0 00h01m29s 32/0/86683
10.0.0.12          65206      0    0    0   0 00h01m29s Active
10.0.0.13          65206      5  10545   0  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 74 Show BGP Summary Field Descriptions

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured local AS setting, if not configured it is the same as the system AS.
BGP Admin State	Down — BGP is administratively disabled. Up — BGP is administratively enabled.
BGP Oper State	Down — BGP is operationally disabled. Up — BGP is operationally enabled.
Confederation AS	The configured confederation AS.
Member Confederations	The configured members of the BGP confederation.

Table 74 Show BGP Summary Field Descriptions (Continued)

Label	Description
Number of Peer Groups	The total number of configured BGP peer groups.
Number of Peers	The total number of configured BGP peers.
Total BGP Active Routes	The total number of BGP routes used in the forwarding table.
Total BGP Routes	The total number of BGP routes learned from BGP peers.
Total BGP Paths	The total number of unique sets of BGP path attributes learned from BGP peers.
Total Path Memory	Total amount of memory used to store the path attributes.
Total Suppressed Routes	Total number of suppressed routes due to route damping.
Total History Routes	Total number of routes with history due to route damping.
Total Decayed Routes	Total number of decayed routes due to route damping.
Neighbor	BGP neighbor address.
AS (Neighbor)	BGP neighbor autonomous system number.
PktRcvd	Total number of packets received from the BGP neighbor.
PktSent	Total number of packets sent to the BGP neighbor.
InQ	The number of BGP messages to be processed.
OutQ	The number of BGP messages to be transmitted.
Up/Down	The amount of time that the BGP neighbor has either been established or not established depending on its current state.
State Recv/Actv/Sent	The BGP neighbor's current state (if not established) or the number of received routes, active routes and sent routes (if established).

ecmp

Syntax `ecmp`

Context `show>router`

Description This command displays the ECMP settings for the router.

Output The following output is an example of router ECMP settings, and [Table 75](#) describes the output fields.

Sample Output

```
*A:ALA-12# show router 3 ecmp
=====
Router ECMP
=====
Instance      Router Name      ECMP      Configured-ECMP-Routes
-----
1             Base             True      8
=====
*A:ALA-12#
```

Table 75 Show ECMP Settings Field Descriptions

Label	Description
Instance	The router instance number.
Router Name	The name of the router instance.
ECMP	False — ECMP is disabled for the instance. True — ECMP is enabled for the instance.
Configured-ECMP-Routes	The number of ECMP routes configured for path sharing.

interface

Syntax **interface** *[[ip-address | ip-int-name] [detail]]* | *[summary]* | *[exclude-services]*

Context show>router

Description This command displays the router IP interface table sorted by interface index.

Parameters *ip-address* — Only displays the interface information associated with the specified IP address.

Values ipv4-prefix: a.b.c.d
 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

ip-int-name — Only displays the interface information associated with the specified IP interface name. 32 characters maximum.

detail — Displays detailed IP interface information.

summary — Displays summary IP interface information for the router.

exclude-services — Displays IP interface information, excluding IP interfaces configured for customer services. Only core network IP interfaces are displayed.

Output The following output is an example of standard IP interface information, and [Table 76](#) describes the output fields.

Sample Output

```
*A:ALA-12# show router 3 interface
=====
Interface Table
=====
Interface-Name      Type  IP-Address      Adm   Opr   Mode
-----
system              Pri   10.10.0.3/32    Up    Up    Network
to-ser1             Pri   10.10.13.3/24   Up    Up    Network
to-ser4             Pri   10.10.34.3/24   Up    Up    Network
to-ser5             Pri   10.10.35.3/24   Up    Up    Network
to-ser6             n/a   n/a             Up    Down  Network
to-web              Pri   10.1.1.3/24     Up    Down  Service
management          Pri   192.168.2.93/20 Up    Up    Network
=====
*A:ALA-12#
```

```
*A:ALA-12# show router 3 interface 10.10.0.3/32
=====
Interface Table
=====
Interface-Name      Type  IP-Address      Adm   Opr   Mode
-----
system              Pri   10.10.0.3/32    Up    Up    Network
=====
SR4#
```

```
*A:ALA-12# show router 3 interface to-ser1
=====
Interface Table
=====
Interface-Name      Type  IP-Address      Adm   Opr   Mode
-----
to-ser1             Pri   10.10.13.3/24   Up    Up    Network
=====
*A:ALA-12#
```

```
*A:ALA-12# show router 3 interface exclude-services
=====
Interface Table
=====
Interface-Name      Type  IP-Address      Adm   Opr   Mode
-----
system              Pri   10.10.0.3/32    Up    Up    Network
to-ser1             Pri   10.10.13.3/24   Up    Up    Network
to-ser4             Pri   10.10.34.3/24   Up    Up    Network
=====
```

```

to-ser5                Pri  10.10.35.3/24      Up    Up    Network
to-ser6                n/a   n/a              Up    Down  Network
management            Pri  192.168.2.93/20      Up    Up    Network
=====
*A:ALA-12#

```

Table 76 Standard IP Interface Field Descriptions

Label	Description
Interface-Name	The IP interface name.
Type	n/a — No IP address has been assigned to the IP interface, so the IP address type is not applicable. Pri — The IP address for the IP interface is the Primary address on the IP interface. Sec — The IP address for the IP interface is a secondary address on the IP interface.
IP-Address	The IP address and subnet mask length of the IP interface. n/a — Indicates no IP address has been assigned to the IP interface.
Adm	Down — The IP interface is administratively disabled. Up — The IP interface is administratively enabled.
Opr	Down — The IP interface is operationally disabled. Up — The IP interface is operationally enabled.
Mode	Network — The IP interface is a network/core IP interface. Service — The IP interface is a service IP interface.

Detailed IP Interface Output

The following output is an example of detailed IP interface information, and [Table 77](#) describes output fields.

Sample Output (7450 ESS and 7750 SR)

```

*A:ALA-12# show router 3 interface detail
=====
Interface Table
=====
Interface
-----
If Name       : to-ser1
Admin State   : Up                               Oper State    : Up

IP Addr/mask  : 10.10.13.3/24                    Address Type   : Primary
IGP Inhibit   : Disabled                          Broadcast Address: Host-ones

```

```

IP Addr/mask : 10.200.0.1/16
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
IP MTU        : 1500
Cflowd       : none

If Type       : Network
Ingress Filter : 100
SNTP Broadcast : False
Arp Timeout   : 14400
ICMP Mask Reply : True

ICMP Details
Redirects      : Disabled
Unreachables  : Number - 100
TTL Expired   : Number - 100
Time (seconds) - 10
Time (seconds) - 10
=====
*A:ALA-12#

```

Table 77 Detailed IP Interface Field Descriptions

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 will be recognized as a local interface by the IGP. Enabled — The secondary IP address on the interface will not be 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.

Table 77 Detailed IP Interface Field Descriptions (Continued)

Label	Description (Continued)
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.
IP MTU	The IP Maximum Transmission Unit (MTU) for the IP interface.
ICMP Mask Reply	False — The IP interface will not reply to a received ICMP mask request. True — The IP interface will reply to a received ICMP mask request.
Cflowd	Specifies the type of Cflowd analysis that is applied to the interface. acl — ACL Cflowd analysis is applied to the interface. interface — Interface cflowd analysis is applied to the interface. none — No Cflowd analysis is applied to the interface.
Redirects	Specifies the maximum number of ICMP redirect messages the IP interface will issue in a given period of time (time in s). Disabled — Indicates the IP interface will not generate ICMP redirect messages.
Unreachables	Specifies the maximum number of ICMP destination unreachable messages the IP interface will issue in a given period of time. Disabled — Indicates the IP interface will not generate ICMP destination unreachable messages.
TTL Expired	The maximum number (Number) of ICMP TTL expired messages the IP interface will issue in a given period of time (Time (seconds)). Disabled — Indicates the IP interface will not generate ICMP TTL expired messages.

Summary IP Interface Output

The following output is an example of summary IP interface information, and [Table 78](#) describes the output fields.

Sample Output

```
*A:ALA-12# show router 3 interface summary
=====
Router Summary (Interfaces)
=====
Instance  Router Name                Interfaces  Admin-Up  Oper-Up
-----
1          Base                      7          7          5
=====
*A:ALA-12#
```

Table 78 Summary IP Interface Field Descriptions

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.
Admin-Up	The number of administratively enabled IP interfaces in the router instance.
Oper-Up	The number of operationally enabled IP interfaces in the router instance.

ldp

Syntax ldp

Context show>router

Description This command displays LDP information.

bindings

Syntax bindings

Context show>router>ldp

Description This command displays LDP bindings information.

Output The following output is an example of LDP bindings information.

Sample Output

```
*A:Dut-A# show router ldp bindings active
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
```

```

(S) - Static          (M) - Multi-homed Secondary Support
(B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
=====
LDP IPv4 Prefix Bindings (Active)
=====
Prefix                Op    IngLbl    EgrLbl    EgrIntf/LspId  EgrNextHop
-----
10.20.1.1/32          Pop  131071    --        --              --
10.20.1.2/32          Push --        131071    1/1/1          10.10.1.2
10.20.1.2/32          Swap 131070    131071    1/1/1          10.10.1.2
10.20.1.2/32          Push --        262141BU  1/1/2          10.10.2.3
10.20.1.2/32          Swap 131070    262141BU  1/1/2          10.10.2.3
10.20.1.3/32          Push --        131069BU  1/1/1          10.10.1.2
10.20.1.3/32          Swap 131069    131069BU  1/1/1          10.10.1.2
10.20.1.3/32          Push --        262143    1/1/2          10.10.2.3
10.20.1.3/32          Swap 131069    262143    1/1/2          10.10.2.3
10.20.1.4/32          Push --        131068    1/1/1          10.10.1.2
10.20.1.4/32          Swap 131068    131068    1/1/1          10.10.1.2
10.20.1.4/32          Push --        262140BU  1/1/2          10.10.2.3
10.20.1.4/32          Swap 131068    262140BU  1/1/2          10.10.2.3
10.20.1.5/32          Push --        131067BU  1/1/1          10.10.1.2
10.20.1.5/32          Swap 131067    131067BU  1/1/1          10.10.1.2
10.20.1.5/32          Push --        262139    1/1/2          10.10.2.3
10.20.1.5/32          Swap 131067    262139    1/1/2          10.10.2.3
10.20.1.6/32          Push --        131066    1/1/1          10.10.1.2
10.20.1.6/32          Swap 131066    131066    1/1/1          10.10.1.2
10.20.1.6/32          Push --        262138BU  1/1/2          10.10.2.3
10.20.1.6/32          Swap 131066    262138BU  1/1/2          10.10.2.3
-----

```

```

-----
No. of IPv4 Prefix Active Bindings: 10
=====

```

```

=====
LDP IPv6 Prefix Bindings (Active)
=====

```

```

Prefix                Op    IngLbl    EgrLbl
EgrNextHop            EgrIf/LspId
-----

```

```

No Matching Entries Found
=====

```

```

=====
LDP Generic IPv4 P2MP Bindings (Active)
=====

```

```

P2MP-Id                Interface
RootAddr              Op    IngLbl    EgrLbl
EgrNH                  EgrIf/LspId
-----

```

```

No Matching Entries Found
=====

```

```

=====
LDP Generic IPv6 P2MP Bindings (Active)
=====

```

```

P2MP-Id                Interface
RootAddr              Op    IngLbl    EgrLbl
EgrNH                  EgrIf/LspId
-----

```



```

-----
No Matching Entries Found
=====

=====
LDP In-Band-SSM IPv4 P2MP Bindings (Active)
=====
Source
Group                                Interface
RootAddr                            Op          IngLbl    EgrLbl
EgrNH                               EgrIf/LspId
-----
No Matching Entries Found
=====

=====
LDP In-Band-SSM IPv6 P2MP Bindings (Active)
=====
Source
Group                                Interface
RootAddr                            Op          IngLbl    EgrLbl
EgrNH                               EgrIf/LspId
-----
No Matching Entries Found
=====

=====
LDP In-Band-VPN-SSM IPv4 P2MP Bindings (Active)
=====
Source
Group                                RD          Op
RootAddr                            Interface   IngLbl    EgrLbl
EgrNH                               EgrIf/LspId
-----
No Matching Entries Found
=====

=====
LDP In-Band-VPN-SSM IPv6 P2MP Bindings (Active)
=====
Source
Group                                RD          Op
RootAddr                            Interface   IngLbl    EgrLbl
EgrNH                               EgrIf/LspId
-----
No Matching Entries Found
=====

*A:Dut-A# show router ldp bindings

=====
LDP Bindings (IPv4 LSR ID 10.1.1.1:0)
              (IPv6 LSR ID ::[0])
=====
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, WP - Label Withdraw Pending, C - Cpipe Service
BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value)

=====

LDP IPv4 Prefix Bindings

=====

Prefix	Peer	IngLbl	EgrLbl	EgrIntf/ LspId	EgrNextHop
10.20.1.1/32	10.20.1.2	131071U	--	--	--
10.20.1.1/32	10.20.1.3	131071U	--	--	--
10.20.1.2/32	10.20.1.2	--	131071	1/1/1	10.10.1.2
10.20.1.2/32	10.20.1.3	131070U	262141	1/1/2	10.10.2.3
10.20.1.3/32	10.20.1.2	131069U	131069	1/1/1	10.10.1.2
10.20.1.3/32	10.20.1.3	--	262143	1/1/2	10.10.2.3
10.20.1.4/32	10.20.1.2	131068N	131068	1/1/1	10.10.1.2
10.20.1.4/32	10.20.1.3	131068BU	262140	1/1/2	10.10.2.3
10.20.1.5/32	10.20.1.2	131067U	131067	1/1/1	10.10.1.2
10.20.1.5/32	10.20.1.3	131067N	262139	1/1/2	10.10.2.3
10.20.1.6/32	10.20.1.2	131066N	131066	1/1/1	10.10.1.2
10.20.1.6/32	10.20.1.3	131066BU	262138	1/1/2	10.10.2.3

No. of IPv4 Prefix Bindings: 12

=====

=====

LDP IPv6 Prefix Bindings

=====

Prefix	IngLbl	EgrLbl
Peer	EgrIntf/LspId	
EgrNextHop		

No Matching Entries Found

=====

=====

LDP Generic IPv4 P2MP Bindings

=====

P2MP-Id	Interface	IngLbl	EgrLbl
RootAddr	EgrIf/LspId		
EgrNH			
Peer			
100			
10.1.1.1	Unknw	--	131051
10.90.90.2	1/1/6		
10.2.2.2:0			
104			
10.1.1.1	Unknw	--	131050
10.90.90.2	1/1/6		
10.2.2.2:0			
600			
10.1.1.1	Unknw	--	131049
10.90.90.2	1/1/6		
10.2.2.2:0			
700			
10.1.1.1	Unknw	--	131048

10.90.90.2	1/1/6		
10.2.2.2:0			
800			
10.1.1.1	Unknw	--	131047
10.90.90.2	1/1/6		
10.2.2.2:0			
900			
10.1.1.1	Unknw	--	131046
10.90.90.2	1/1/6		
10.2.2.2:0			
1500			
10.1.1.1	Unknw	--	131045
10.90.90.2	1/1/6		
10.2.2.2:0			
100			
10.6.6.6	Unknw	--	131044
10.90.90.2	1/1/6		
10.2.2.2:0			
900			
10.6.6.6	Unknw	--	131043
10.90.90.2	1/1/6		
10.2.2.2:0			

No. of Generic IPv4 P2MP Bindings: 9
=====

=====

LDP Generic IPv6 P2MP Bindings
=====

P2MP-Id	Interface	IngLbl	EgrLbl
RootAddr	EgrIf/LspId		
EgrNH			
Peer			

No Matching Entries Found
=====

=====

LDP In-Band-SSM IPv4 P2MP Bindings
=====

Source	Interface	IngLbl	EgrLbl
Group	EgrIf/LspId		
RootAddr			
EgrNH			
Peer			

No Matching Entries Found
=====

=====

LDP In-Band-SSM IPv6 P2MP Bindings
=====

Source

```

Group
RootAddr
EgrNH
Peer
-----
No Matching Entries Found
=====

=====
LDP In-Band-VPN-SSM IPv4 P2MP Bindings
=====
Source
Group
RootAddr
EgrNH
Peer
-----
RD
Interface
EgrIf/LspId
IngLbl
EgrLbl
-----
10.1.1.1
225.0.0.1
10.3.3.3
10.60.60.1
10.2.2.2:100
10.1.1.1:100
Unknwn
1/1/1
--
100

10.1.1.1
225.0.0.1
10.3.3.3
10.60.60.1
10.2.2.2:100
10.1.1.1:100
Unknwn
1/1/1
--
100

10.1.1.1
225.0.0.1
10.3.3.3
10.60.60.1
10.2.2.2:100
10.1.1.1:100
Unknwn
1/1/1
--
100

-----
No. of In-Band-VPN-SSM IPv4 P2MP Bindings: 3
=====

=====
LDP In-Band-VPN-SSM IPv6 P2MP Bindings
=====
Source
Group
RootAddr
EgrNH
Peer
-----
RD
Interface
EgrIf/LspId
IngLbl
EgrLbl
-----
10.1.1.1
225.0.0.1
2000::3000
10.60.60.1
10.2.2.2:100
10.1.1.1:100
Unknwn
1/1/1
--
100

10.1.1.1
225.0.0.1
2000::3000
10.60.60.1
10.2.2.2:100
10.1.1.1:100
Unknwn
1/1/1
--
100

```

10.2.2.2:100

10.1.1.1

225.0.0.1

2000::3000

10.60.60.1

10.2.2.2:100

10.1.1.1:100

Unknwn

--

100

1/1/1

No. of In-Band-VPN-SSM IPv6 P2MP Bindings: 3
=====

=====

LDP Service FEC 128 Bindings

Type	VCId	SDPID	IngLbl	LMTU
Peer	SvcId		EgrLbl	RMTU
?-Eth	100	R. Src	--	None
10.2.2.2:0	Ukwn		131023D	986
?-Eth	500	R. Src	--	None
10.2.2.2:0	Ukwn		131022D	1386
?-Eth	2001	R. Src	--	None
10.2.2.2:0	Ukwn		131019D	986
?-Eth	2003	R. Src	--	None
10.2.2.2:0	Ukwn		131017D	986
?-Ipipe	1800	R. Src	--	None
10.2.2.2:0	Ukwn		131014D	1486

No. of VC Labels: 5
=====

=====

LDP Service FEC 129 Bindings

SAII	AGII	IngLbl	LMTU
TAII	Type	EgrLbl	RMTU
Peer	SvcId	SDPID	

No Matching Entries Found
=====

mvpn

Syntax mvpn

Context show>router router-instance

Description This command displays Multicast VPN related information. The router instance must be specified.

Output The following output is an example of MVPN information.

Sample Output

```
*A:Dut-C# show router 1 mvpn
=====
MVPN 1 configuration data
=====
signaling          : Bgp                auto-discovery      : Enabled
UMH Selection      : Highest-Ip          intersite-shared     : Enabled
vrf-import         : N/A
vrf-export         : N/A
vrf-target         : target:1:1
C-Mcast Import RT  : target:10.20.1.3:2

ipmsi              : pim-asm 239.1.1.1
admin status       : Up                  three-way-hello      : N/A
hello-interval     : N/A                  hello-multiplier     : 35 * 0.1
tracking support    : Disabled             Improved Assert      : N/A

spmsi              : pim-ssm 225.0.0.0/32
join-tlv-packing   : N/A
data-delay-interval: 3 seconds
data-threshold     : 224.0.0.0/4 --> 1 kbps
=====
```

rip

Syntax **rip**

Context show>router

Description Displays RIP information.

database

Syntax **database** [*ip-address* [*/mask*] [*longer*]] [*peer ip-address*] [*detail* [*qos*]]

Context show>router>rip

Description Displays all routes in the RIP database.

Parameters *ip-address /mask* — Specifies the IP address.

Values ip-address: a.b.c.d.
mask: 1 to 32

longer — Specifies the longer prefix match entries should also be displayed.

peer ip-address — Specifies the peer IP address.

Values a.b.c.d

detail — Displays detailed information.

qos — Specifies the QoS.

Output The following output is an example of RIP route database information, and [Table 79](#) describes the output fields.

Sample Output

```
*A:ALA-1# show rip database
=====
RIP Route Database
=====
Destination      Peer      NextHop      Metric  Tag      TTL      Valid
-----
10.0.0.10/32     10.1.7.15  0.0.0.0      2       0x0000   163     No
10.0.0.10/32     10.1.8.14  0.0.0.0      2       0x0000   179     No
10.0.0.14/32     10.1.8.14  0.0.0.0      1       0x0000   179     Yes
10.0.6.0/24      10.1.7.15  0.0.0.0      11      0x2002   163     No
10.0.6.0/24      10.1.8.14  0.0.0.0      11      0x2002   179     No
10.0.7.0/24      10.1.7.15  0.0.0.0      11      0x2002   163     No
10.0.7.0/24      10.1.8.14  0.0.0.0      11      0x2002   179     No
10.1.5.0/24      10.1.7.15  0.0.0.0      2       0x0000   151     Yes
10.1.5.0/24      10.1.8.14  0.0.0.0      1       0x0000   167     No
10.100.17.16/31  10.1.7.15  0.0.0.0      2       0x0000   151     No
10.100.17.16/31  10.1.8.14  0.0.0.0      2       0x0000   167     No
-----
No. of Routes: 11
=====
*A:ALA-12#
```

Table 79 Show RIP Database Field Descriptions

Label	Description
Destination	The RIP destination for the route.
Peer	The router ID of the peer router.
NextHop	The IP address of the next hop.
Metric	The hop count to rate the value of different hops.
Tag	The value to distinguish between internal routes (learned by RIP) and external routes (learned from other protocols).
TTL	Displays how many seconds the specific route will remain in the routing table. When an entry reaches 0, it is removed from the routing table.
Valid	No — The route is not valid. Yes — The route is valid.

neighbor

- Syntax** `neighbor [ip-address | ip-int-name] [detail] [advertised-routes]`
- Context** `show>router>rip`
- Description** Displays RIP neighbor interface information.
- Parameters** *ip-address* | *ip-int-name* — Displays information for the specified IP interface.
- Values** *i*[-address: a.b.c.d
ip-int-name: 32 chars max
- Default** all neighbor interfaces
- advertised-routes** — Displays the routes advertised to RIP neighbors. If no neighbors are specified, then all routes advertised to all neighbors are displayed. If a specific neighbor is given then only routes advertised to the given neighbor/interface are displayed.
- Default** displays RIP information
- Output** The following output is an example of standard RIP group information, and [Table 80](#) describes the output fields.

Sample Output

```
*A:ALA-12# show router 3 rip neighbor
=====
RIP Neighbors
=====
Interface                               Adm  Opr  Primary IP      Send  Recv  Metric
                                   Mode  Mode  In
-----
router-21/1                            Up    Up    10.0.3.12       None  Both  1
router-21/2                            Up    Up    10.0.5.12       BCast Both  1
router-21/3                            Up    Up    10.0.6.12       BCast Both  1
router-21/4                            Up    Up    10.0.10.12      BCast Both  1
router-21/5                            Up    Up    10.0.9.12       BCast Both  1
router-21/6                            Up    Up    10.0.17.12      None  Both  1
router-21/7                            Up    Up    10.0.16.12      None  Both  1
=====
*A:ALA-12#
```

Table 80 Standard Show RIP Neighbor Field Descriptions

Label	Description
Neighbor	The RIP neighbor interface name.
Adm	Down — The RIP neighbor interface is administratively down. Up — The RIP neighbor interface is administratively up.

Table 80 Standard Show RIP Neighbor Field Descriptions (Continued)

Label	Description
Opr	Down — The RIP neighbor interface is operationally down. Up — The RIP neighbor interface is operationally up.
Primary IP	The primary IP address of the RIP neighbor interface.
Send Mode	Bcast — Specifies that RIPv2 formatted messages are sent to the broadcast address. Mcast — Specifies that RIPv2 formatted messages are sent to the multicast address. None — Specifies that no RIP messages are sent (i.e., silent listener). RIPv1 — Specifies that RIPv1 formatted messages are sent to the broadcast address.
Recv Mode	Both — Specifies that RIP updates in either version 1 or version 2 format will be accepted. None — Specifies that RIP updates will not be accepted. RIPv1 — Specifies that RIP updates in version 1 format only will be accepted. RIPv2 — Specifies that RIP updates in version 2 format only will be accepted.
Metric In	The metric added to routes received from a RIP neighbor.

Detailed Show RIP Neighbor Output

The following output is an example of detailed RIP group information, and [Table 81](#) describes the output fields.

Sample Output

```
*A:ALA-12# show router 3 rip peers
=====
RIP Peers
=====
Peer IP Addr      Interface Name      Version      Last Update
-----
10.0.5.13         router-2/2          RIPv2        0
10.0.6.16         router-2/3          RIPv2        2
10.0.9.14         router-2/5          RIPv2        8
10.0.10.15        router-2/4          RIPv2        0
-----
No. of Peers: 4
=====
*A:ALA-12#

*A:ALA-12# show router 3 rip neighbor detail
```

```

=====
RIP Neighbors (Detail)
=====
Neighbor "router-2/7"
-----
Description      : No Description Available
Primary IP       : 10.0.16.12          Group          : seven
Admin State      : Up                  Oper State      : Up
Send Mode        : None                Receive Mode    : Both
Metric In        : 1                   Metric Out      : 1
Split Horizon    : Enabled             Check Zero      : Disabled
Message Size     : 25                  Preference     : 100
Auth. Type       : None                Update Timer    : 3
Timeout Timer    : 6                   Flush Timer     : 6
Export Policies:
    Rip2Rip
    direct2Rip
    bgp2Rip
Import Policies:
    None
=====
*A:ALA-12#

```

Table 81 Detailed Show RIP Neighbor Field Descriptions

Label	Description
Neighbor	The RIP neighbor name.
Description	The RIP neighbor description. No Description Available indicates no description is configured.
Primary IP	The RIP neighbor interface primary IP address.
Group	The RIP group name of the neighbor interface.
Admin State	Down — The RIP neighbor interface is administratively down. Up — The RIP neighbor interface is administratively up.
Oper State	Down — The RIP neighbor interface is operationally down. Up — The RIP neighbor interface is operationally up.
Send Mode	Bcast — Specifies that RIPv2 formatted messages are sent to the broadcast address. Mcast — Specifies that RIPv2 formatted messages are sent to the multicast address. None — Specifies that no RIP messages are sent (i.e., silent listener). RIPv1 — Specifies that RIPv1 formatted messages are sent to the broadcast address.
Recv Mode	Both — Specifies that RIP updates in either version 1 or version 2 format will be accepted. None — Specifies that RIP updates will not be accepted. RIPv1 — Specifies that RIP updates in version 1 format only will be accepted. RIPv2 — Specifies that RIP updates in version 2 format only will be accepted.
Metric In	The metric value added to routes received from a RIP neighbor.

Table 81 Detailed Show RIP Neighbor Field Descriptions (Continued)

Label	Description (Continued)
Metric Out	The value added to routes exported into RIP and advertised to RIP neighbors.
Split Horizon	Disabled — Split horizon disabled for the neighbor. Enabled — Split horizon and poison reverse enabled for the neighbor.
Check Zero	Disabled — Checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications are not checked allowing receipt of RIP messages even if mandatory zero fields are non-zero for the neighbor. Enabled — Checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications and rejecting non-compliant RIP messages is enabled for the neighbor.
Message Size	The maximum number of routes per RIP update message.
Preference	The preference of RIP routes from the neighbor.
Auth. Type	Specifies the authentication type.
Update Timer	The current setting of the RIP update timer value expressed in seconds.
Timeout Timer	The current RIP timeout timer value expressed in seconds.
Export Policies	The export route policy that is used to determine routes advertised to all peers.
Import Policies	The import route policy that is used to determine which routes are accepted from RIP neighbors.

Sample Output

```
*A:ALA-12# show router 3 rip neighbors interface advertised-routes
=====
RIP Advertised Routes
=====
Destination      Interface      NextHop      Metric  Tag      TTL
-----
10.0.0.2/32      10.1.8.12     0.0.0.0      10      0x2002   n/a
10.0.0.5/32      10.1.8.12     0.0.0.0      10      0x2002   n/a
10.0.0.8/32      10.1.8.12     0.0.0.0      10      0x2002   n/a
10.0.0.9/32      10.1.8.12     0.0.0.0      10      0x2002   n/a
10.0.0.10/32     10.1.8.12     0.0.0.0      10      0x2002   n/a
10.0.0.12/32     10.1.8.12     0.0.0.0      1       0x0000   n/a
10.0.0.13/32     10.1.8.12     0.0.0.0      10      0x2002   n/a
10.0.0.14/32     10.1.8.12     0.0.0.0      16      0x0000   n/a
10.0.0.15/32     10.1.8.12     0.0.0.0      2       0x0000   n/a
10.0.0.16/32     10.1.8.12     0.0.0.0      3       0x0000   n/a
-----
No. of Advertised Routes: 10
=====
*A:ALA-12#
```

peer

Syntax	peer [<i>interface-name</i>]
Context	show>router>rip
Description	Displays RIP peer information.
Parameters	<i>ip-int-name</i> — Displays peer information for peers on the specified IP interface. 32 characters maximum. Default Display peers for all interfaces.
Output	Show RIP Peer Output

[Table 82](#) describes the command output fields for a RIP peer:

Table 82 Show RIP Peer Field Descriptions

Label	Description
Peer IP Addr	The IP address of the peer router.
Interface Name	The peer interface name.
Version	The version of RIP running on the peer.
Last Update	The number of days since the last update.
No. of Peers	The number of RIP peers.

statistics

Syntax	statistics [<i>ip-addr</i> <i>ip-int-name</i>]
Context	show>router>rip
Description	Display Interface level statistics for the RIP protocol. If no IP address or interface name is specified, then all configured RIP interfaces are displayed. If an IP address or interface name is specified, then only data regarding the specified RIP interface is displayed.
Parameters	<i>ip-addr</i> <i>ip-int-name</i> — Displays statistics for the specified IP interface. Values ip-int-name: 32 chars max

ip-address: a.b.c.d

Output The following output is an example of RIP statistics information, and [Table 83](#) describes the output fields.

Sample Output

```
RIP Statistics
=====
Learned Routes      : 0                Timed Out Routes   : 0
Current Memory      : 120624           Maximum Memory     : 262144

-----
Interface "to-web"
-----
Primary IP          : 10.1.1.3          Update Timer        : 30
Timeout Timer       : 180              Flush Timer         : 120
Counter              Total              Last 5 Min          Last 1 Min
-----
Updates Sent         0                  0                  0
Triggered Updates    0                  0                  0
Bad Packets Received 0                  0                  0
RIPv1 Updates Received 0              0                  0
RIPv1 Updates Ignored 0              0                  0
RIPv1 Bad Routes     0                  0                  0
RIPv1 Requests Received 0            0                  0
RIPv1 Requests Ignored 0            0                  0
RIPv2 Updates Received 0              0                  0
RIPv2 Updates Ignored 0              0                  0
RIPv2 Bad Routes     0                  0                  0
RIPv2 Requests Received 0            0                  0
RIPv2 Requests Ignored 0            0                  0
Authentication Errors 0                  0                  0
=====
*A:ALA-12#
```

Table 83 Show RIP Statistics Field Descriptions

Label	Description
Learned Routes	The number of RIP-learned routes were exported to RIP neighbors.
Timed Out Routes	The number of routes that have been timed out.
Current Memory	The amount of memory used by this RIP router instance.
Maximum Memory	The amount of memory allocated for this RIP router instance.
Interface	Displays the name of each interface configured in RIP and associated RIP statistics.
Primary IP	The interface IP address.
Update Timer	The current setting of the RIP update timer value expressed in seconds.
Timeout Timer	The current RIP timeout timer value expressed in seconds.

Table 83 Show RIP Statistics Field Descriptions (Continued)

Label	Description
Flush Timer	The number of seconds after a route has been declared invalid that it is flushed from the route database.
Updates Sent	Total — The total number of RIP updates that were sent. Last 5 Min — The number of RIP updates that were sent in the last 5 minutes. Last 1 Min — The number of RIP updates that were sent in the last 1 minute.
Triggered Updates	Total — The total number of triggered updates sent. These updates are sent before the entire RIP routing table is sent. Last 5 Min — The number of triggered updates that were sent in the last 5 minutes. Last 1 Min — The number of triggered updates that were sent in the last 1 minute.
Bad Packets Received	Total — The total number of RIP updates received on this interface that were discarded as invalid. Last 5 Min — The number of RIP updates received on this interface that were discarded as invalid in the last 5 minutes. Last 1 Min — The number of RIP updates received on this interface that were discarded as invalid in the last 1 minute.
RIPv1 Updates Received	Total — The total number of RIPv1 updates received. Last 5 Min — The number of RIPv1 updates received in the last 5 minutes. Last 1 Min — The number of RIPv1 updates received in the last 1 minute.
RIPv1 Updates Ignored	Total — The total number of RIPv1 updates ignored. Last 5 Min — The number of RIPv1 updates ignored in the last 5 minutes. Last 1 Min — The number of RIPv1 updates ignored in the last 1 minute.
RIPv1 Bad Routes	Total — The total number of bad routes received from the peer. Last 5 Min — The number of bad routes received from the peer in the last 5 minutes. Last 1 Min — The number of bad routes received from the peer in the last minute.
RIPv1 Requests Received	Total — The total number of times the router received RIPv1 route requests from other routers. Last 5 Min — The number of times the router received RIPv1 route requests from other routers in the last 5 minutes. Last 1 Min — The number of times the router received RIPv1 route requests from other routers in the last 1 minute.
RIPv1 Requests Ignored	Total — The total number of times the router ignored RIPv1 route requests from other routers. Last 5 Min — The number of times the router ignored RIPv1 route requests from other routers in the last 5 minutes. Last 1 Min — The number of times the router ignored RIPv1 route requests from other routers in the last 1 minute.

Table 83 Show RIP Statistics Field Descriptions (Continued)

Label	Description
RIPv2 Updates Received	Total — The total number of RIPv2 updates received. Last 5 Min — The number of RIPv2 updates received in the last 5 minutes. Last 1 Min — The number of RIPv2 updates received in the last minute.
RIPv2 Updates Ignored	Total — The total number of RIPv2 updates ignored. Last 5 Min — The number of RIPv2 updates ignored in the last 5 minutes. Last 1 Min — The number of RIPv2 updates ignored in the last minute.
RIPv2 Bad Routes	Total — The total number of bad routes received from the peer. Last 5 Min — The number of bad routes received from the peer in the last 5 minutes. Last 1 Min — The number of bad routes received from the peer in the last minute.
RIPv2 Requests Received	Total — The total number of times the router received RIPv2 route requests from other routers. Last 5 Min — The number of times the router received RIPv2 route requests from other routers in the last 5 minutes. Last 1 Min — The number of times the router received RIPv2 route requests from other routers in the last minute.
RIPv2 Requests Ignored	Total — The total number of times the router ignored RIPv2 route requests from other routers. Last 5 Min — The number of times the router ignored RIPv2 route requests from other routers in the last 5 minutes. Last 1 Min — The number of times the router ignored RIPv2 route requests from other routers in the last minute.
Authentication Errors	Total — The total number of authentication errors to secure table updates. Last 5 Min — The number of authentication errors to secure table updates in the last 5 minutes. Last 1 Min — The number of authentication errors to secure table updates in the last minute.

route-table

Syntax `route-table [family] [ip-prefix/prefix-length] [longer | exact | protocol protocol-name] [all] [next-hop-type type] [qos] [alternative] [accounting-class]`
`route-table [family] [summary]`
`route-table tunnel-endpoints [ip-prefix/prefix-length] [longer | exact] [detail]`
`route-table [family] [ip-prefix/prefix-length] [longer | exact | protocol protocol-name] extensive [all]`

Context show>router

Description This command displays the active routes in the routing table.

If no command line arguments are specified, all routes are displayed, sorted by prefix.

Parameters

family — Specifies the family.

Values ipv4, ipv6, mcast-ipv4, mcast-ipv6

ip-prefix[/prefix-length] — Displays routes only matching the specified *ip-prefix* and optional *mask*.

Values ipv4-prefix: a.b.c.d (host bits must be 0)

ipv4-prefix-le: 0 to 32

ipv6-prefix:

- x:x:x:x:x:x:x (eight 16-bit pieces)

- x:x:x:x:x:d.d.d.d

- x: [0 to FFFF] H

- d: [0 to 255] D

ipv6-prefix-le: 0 to 128

longer — Displays routes matching the *ip-prefix/mask* and routes with longer masks.

exact — Displays the exact route matching the *ip-prefix/mask* masks.

protocol-name — Displays routes learned from the specified protocol.

The following values apply to the 7750 SR and the 7450 ESS:

Values local, host, sub-mgmt, managed, static, ldp, ospf, ospf3, isis, rip, bgp, bgp-vpn, bgp-evpn, aggregate, vpn-leak, tms, nat, periodic, dhcpv6-pd, dhcpv6-na, dhcpv6-ta, dhcpv6-pd-excl, ripng, ipsec, bgp-label

The following values apply to the 7950 XRS:

Values local, host, managed, static, ldp, ospf, ospf3, isis, rip, bgp, bgp-vpn, bgp-evpn, aggregate, vpn-leak, dhcpv6-pd, dhcpv6-na, dhcpv6-ta, dhcpv6-pd-excl, ripng, bgp-label, ipv4, ipv6, mcast-ipv4, mcast-ipv6

all — Includes inactive routes.

type — Displays only the tunneled next-hops. For each route entry, the tunnel owner and tunnel ID is shown.

Values tunneled

qos — Specifies to display QoS active routes in the routing table.

alternative — Displays LFS and backup route details.

accounting-class — Specifies to display accounting class information.

summary — Displays a route table summary information.

tunnel-endpoints — Specifies to include tunnel endpoint information.

extensive — Displayed detailed information.

Output

The following output is an example of standard route table information, and [Table 84](#) 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 84 Standard Show Route Table Field Descriptions

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.

The following is a sample output of the **show router *router-instance* route-table all** command for a VPRN on the standby PE for prefix 10.13.1.0/24.

Sample Output

```
*A:ALA-12# show router 1 route-table 10.13.1.0/24 all
=====
Route Table (Service: 1)
=====
Dest Prefix[Flags]                                Type      Proto      Age          Pref
      Next Hop[Interface Name]                    Active      Metric
-----
10.13.1.0/24 [E]                                Remote    BGP          00h01m44s    170
      10.3.1.1                                    N          0
10.13.1.0/24                                    Remote    BGP VPN      00h01m19s    170
      10.15.1.100 (tunneled)                      Y          0
=====
```

service-prefix

Syntax	service-prefix
Context	show>router
Description	This command displays service-prefix information.
Output	The following output is an example of service prefix information, and Table 85 describes the output fields.

Sample Output

```
*A:ALA-12# show router 3 service-prefix
=====
Address Ranges Reserved for Services (Service: 3)
=====
IP Prefix          Mask          Exclusive
-----
No Matching Entries Found
=====
*A:ALA-12>show>router#
```

Table 85 Show Service Prefix Field Descriptions

Label	Description
IP Prefix	Displays information for the specified IP prefix.
Mask	Displays information for the specified mask length.

static-arp

Syntax **static-arp** [*ip-address* | *ip-int-name* | *mac ieee-mac-address*]

Context	show>router
Description	This command displays the router static ARP table sorted by IP address. If no options are present, all ARP entries are displayed.
Parameters	<i>ip-address</i> — Only displays static ARP entries associated with the specified IP address. Values a.b.c.d <i>ip-int-name</i> — Only displays static ARP entries associated with the specified IP interface name. 32 characters maximum. <i>ieee-mac-address</i> — Only displays static ARP entries associated with the specified MAC address. Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
Output	The following output is an example of ARP table information, and Table 86 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 86 Static ARP Table Field Descriptions

Label	Description
IP Address	The IP address of the static ARP entry.
MAC Address	The MAC address of the static ARP entry.
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 [<i>family</i>] [<i>ip-prefix/prefix-length</i> preference <i>preference</i> next-hop <i>ip-address</i> tag <i>tag</i>] [detail]
Context	show>router
Description	This command displays the static entries in the routing table. If no options are present, all static routes are displayed sorted by prefix.
Parameters	<i>family</i> — Specifies the family to display. Values ipv4, ipv6, mcast-ipv4, mcast-ipv6 <i>ip-prefix lmask</i> — Displays static routes only matching the specified <i>ip-prefix</i> and <i>mask</i> . Values ipv4-prefix: a.b.c.d (host bits must be 0) ipv4-prefix-le: 0 to 32 ipv6-prefix: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d

- x: [0 to FFFF] H

- d: [0 to 255] D

ipv6-prefix-len: 0 to 128

preference — Only displays static routes with the specified route preference.

Values 0 to 65535

ip-address — Only displays static routes with the specified next hop IP address.

Values ipv4-address: a.b.c.d

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

tag — Displays the tag used to add a 32-bit integer tag to the static route. The tag is used in route policies to control distribution of the route into other protocols.

Values 0 to 4294967295

detail — Displays detailed information about the static route.

Output The following output is an example of static route table information, and [Table 87](#) describes the output fields.

Sample Output

```
*A:ALA-12# show router 3 static-route
=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.250.0/24  5      1      ID   10.200.10.1      to-ser1        Y
192.168.252.0/24  5      1      NH   10.10.0.254      n/a            N
192.168.253.0/24  5      1      NH   to-ser1          n/a            N
192.168.253.0/24  5      1      NH   10.10.0.254      n/a            N
192.168.254.0/24  4      1      BH   black-hole       n/a            Y
=====
*A:ALA-12#

*A:ALA-12# show router 3 static-route 192.168.250.0/24
=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.250.0/24  5      1      ID   10.200.10.1      to-ser1        Y
=====
*A:ALA-12#
```

```
*A:ALA-12# show router 3 static-route preference 4
=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.254.0/24  4      1      BH    black-hole      n/a           Y
=====
*A:ALA-12#
```

```
*A:ALA-12# show router 3 static-route next-hop 10.10.0.254
=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.253.0/24  5      1      NH    10.10.0.254     n/a           N
=====
*A:ALA-12#
```

```
*A:Dut-B# show router static-route

=====
Static Route Table (Router: Base)  Family: IPv4
=====
Prefix      Tag      Met      Pref Type Act
Next Hop    Interface
-----
10.2.3.4/32  0         1        5    NH    Y
10.11.25.6
ip-10.11.25.5_base_to_cpe_static
10.11.15.0/24  0         1        5    NH    Y
10.11.25.6
ip-10.11.25.5_base_to_cpe_static
-----
No. of Static Routes: 2
=====
```

```
*A:Dut-B# show router static-route detail
=====
Static Route Table (Router: Base)  Family: IPv4
=====
Network      : 10.2.3.4/32
Nexthop      : 10.11.25.6
Type         : Nexthop
Interface    : ip-10.11.25.5_base_to_cpe_stat*
Metric       : 1
Admin State  : Up
BFD          : disabled
CPE-check    : enabled
Target       : 10.11.18.6
Interval     : 1
Log          : N
CPE Host Up Time : 0d 00:00:02
CPE Echo Req Tx : 3
CPE Up Trans  : 1
CPE TTL      : 2
Nexthop Type : IP
Active        : Y
Preference    : 5
Tag           : 0
State         : n/a
Drop Count    : 3
CPE Echo Reply Rx : 3
CPE Down Trans : 0
```

```

-----
Network      : 10.11.15.0/24
NextHop      : 10.11.25.6
Type         : Nexthop                               Nexthop Type      : IP
Interface    : ip-10.11.25.5_base_to_cpe_stat*      Active           : Y
Metric       : 1                                     Preference       : 5
Admin State  : Up                                    Tag              : 0
BFD          : disabled
CPE-check    : disabled
-----
No. of Static Routes: 2

=====

*A:CPM133>config>router# show router static-route 10.3.3.3/32 detail

=====
Static Route Table (Router: Base)  Family: IPv4
=====
Prefix      : 10.3.3.3/32
NextHop     : n/a
Type        : Blackhole                               Nexthop Type      : IP
Interface   : n/a                                     Active           : Y
Prefix List : n/a                                     Prefix List Type  : n/a
Metric      : 1                                       Preference       : 5
Admin State : Up                                    Tag              : 0
BFD         : disabled                               Community        : 100:33
CPE-check   : disabled
-----
No. of Static Routes: 1

=====

```

Table 87 **Show Static Route Field Descriptions**

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 black hole route. The Nexthop for this type of route is black-hole. ID — The static route is an indirect route, where the nexthop for this type of route is the non-directly connected next hop. NH — The route is a static route with a directly connected next hop. The Nexthop for this type of route is either the next hop IP address or an egress IP interface name.
Next Hop	The next hop for the static route destination.

Table 87 Show Static Route Field Descriptions (Continued)

Label	Description (Continued)
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 black hole 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 summary [ipv4 ipv6] tunnel-table [protocol protocol] {ipv4 ipv6} tunnel-table [ip-prefix[/mask]] [alternative] [ipv4 ipv6] detail tunnel-table [ip-prefix[/mask]] [alternative] tunnel-table mpls-tp tunnel-table [ip-prefix[/mask]] protocol protocol [detail] tunnel-table [ip-prefix[/mask]] sdp sdp-id
Context	show>router
Description	This command displays tunnel table information. Auto-bind GRE tunnels are not displayed in show command output. GRE tunnels are not the same as SDP tunnels that use the GRE encapsulation type. When the auto-bind-tunnel command is used when configuring a VPRN service, it means the MP-BGP NH resolution is referring to core routing instance for IP reachability. For a VPRN service this object specifies the lookup to be used by the routing instance if no SDP to the destination exists.
Parameters	summary — Displays summary tunnel table information. ip-address[/mask] — Displays the specified tunnel table's destination IP address and mask. Values ipv4-prefix: a.b.c.d (host bits must be 0) ipv4-prefix-le: 0 to 32 ipv6-prefix: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF] H • d: [0 to 255] D

ipv6-prefix-le: 0 to 128

ipv4 — Displays only tunnel table information for IPv4 addresses.

ipv6 — Displays only tunnel table information for IPv6 addresses.

protocol — Displays LDP protocol information.

Values bgp, ldp, rsvp, sdp, ospf, isis, sr-te, fpe, udp

alternative — Displays backup route details.

mpls-tp — Displays MPLS TP tunnel table information.

sdp-id — Displays information pertaining to the specified SDP.

Values 1 to 17407

Output The following output is an example of tunnel table information, and [Table 88](#) describes the output fields.

Sample Output

```
*A:ALA-12>config>service# show router 3 tunnel-table
=====
Tunnel Table
=====
Destination      Owner   Encap   Tunnel Id   Pref      NexthopMetric
-----
10.0.0.1/32      sdp     GRE     10         5         10.0.0.1      0
10.0.0.1/32      sdp     GRE     21         5         10.0.0.1      0
10.0.0.1/32      sdp     GRE     31         5         10.0.0.1      0
10.0.0.1/32      sdp     GRE     41         5         10.0.0.1      0
=====
*A:ALA-12>config>service#

*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 88 Show Tunnel Table Field Descriptions

Label	Description
Destination	The route's destination address and mask.
Owner	Specifies the tunnel owner.
Encap	Specifies the tunnel's encapsulation type.

Table 88 Show Tunnel Table Field Descriptions (Continued)

Label	Description (Continued)
Tunnel ID	Specifies the tunnel (SDP) identifier.
Pref	Specifies the route preference for routes learned from the configured peer(s).
Nexthop	The next hop for the route's destination.
Metric	The route metric value for the route.

statistics

Syntax **statistics** [**interface** *ip-int-name* | *ip-address*]

Context show>router>dhcp

Description This command displays statistics for DHCP Relay and DHCP snooping.

If no IP address or interface name is specified, then all configured interfaces are displayed.

If an IP address or interface name is specified, then only data regarding the specified interface is displayed.

Parameters *ip-int-name* | *ip-address* — Displays statistics for the specified IP interface.

Values ip-int-name: 32 chars max
ip-address: a.b.c.d

Output The following output is an example of DHCP statistics information, and [Table 89](#) describes the output fields.

Sample Output

```
*A:ALA-1# show router dhcp statistics
=====
DHCP Global Statistics
=====
Rx Packets                : 0
Tx Packets                : 0
Rx Malformed Packets     : 0
Rx Untrusted Packets     : 0
Client Packets Discarded  : 0
Client Packets Relayed    : 0
Client Packets Snooped    : 0
Server Packets Discarded  : 0
Server Packets Relayed    : 0
Server Packets Snooped    : 0
=====
*A:ALA-1#
```

Table 89 Show DHCP Statistics Field Descriptions

Label	Description
Received Packets	The number of packets received from the DHCP clients.
Transmitted Packets	The number of packets transmitted to the DHCP clients.
Received Malformed Packets	The number of malformed packets received from the DHCP clients.
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients.
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**

Context show>router>dhcp

Description This command displays the status of the DHCP Relay and DHCP snooping functions on each interface.

Output The following output is an example of DHCP summary information, and [Table 90](#) describes the output fields.

Sample Output

```
A:ALA-48# show router dhcp summary
=====
Interface Name          Arp      Used/    Info    Admin
                        Populate Provided Option State
-----
ies-10-10.10.1.1        Yes      1000/8000 Keep    Up
```

```

ies-100-100.100.1.1      No      0/0      Keep    Down
ies-11-10.11.1.1        Yes     1000/8000 Keep    Up
ies-12-10.12.1.1        Yes     1000/8000 Keep    Up
ies-13-13.13.1.1        Yes     1000/8000 Keep    Up
ies-14-14.14.1.1        Yes     1000/8000 Keep    Up
ies-15-15.15.1.1        Yes     1000/8000 Keep    Up
ies-16-16.16.1.1        No      0/0      Keep    Down
ies-2-10.17.1.1         No      0/0      Keep    Down
ies-8-8.8.1.1           Yes     1000/8000 Keep    Up
ies-9-9.9.1.1           Yes     1000/8000 Keep    Up
-----
Interfaces: 11
=====

```

Table 90 **Show DHCP Summary Field Descriptions**

Label	Description
Interface Name	Name of the router interface.
ARP Populate	Indicates whether or not ARP populate is enabled.
Info Option	Indicates whether Option 82 is enabled.
Admin State	Indicates the administrative status.

wpp

Syntax **wpp**
wpp [**portal** *wpp-portal-name*] [**host** *ip-address*] **hosts**
wpp portal *wpp-portal-name*
wpp statistics

Context show>router

Description This command displays Web Portal Protocol information.

Parameters *wpp-portal-name* — Specifies the name of this WPP portal. 32 chars max
ip-address — Specifies the host IP address.

Values a.b.c.d

hosts — Displays the hosts enabled on the portal.

statistics — Displays WPP statistics.

3.9.2.2 VPRN Clear Commands

arp

Syntax	arp
Context	clear>service>id
Description	This command clears ARP data.

arp-host

Syntax	arp-host { all mac <i>ieee-address</i> sap <i>sap-id</i> ip-address <i>ip-address</i> [/ <i>mask</i>]} arp-host { port <i>port-id</i> { inter-dest-id <i>intermediate-destination-id</i> no-inter-dest-id } [port <i>port-id</i>]} arp-host statistics [sap <i>sap-id</i> interface <i>interface-name</i>]		
Context	clear>service>id		
Description	This command clears ARP host data.		
Parameters	all — Clears all ARP host statistics. <i>ieee-address</i> — Clears the ARP host MAC address information. 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. Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx <i>sap-id</i> — Clears the specified SAP information. <i>ip-address</i> [/ <i>mask</i>] — Clears the specified IP address and mask. Values a.b.c.d mask: 1 to 32 <i>port-id</i> — Clear the specified port ID information. Values <i>port-id</i> <i>slot/mda/port</i> [.<i>channel</i>] <i>bundle-id</i> <i>bundle-type-slot/mda.bundle-num</i> <i>bundle</i> keyword <i>type</i> ppp <i>bundle-num</i> 1 to 336 <i>bpgrp-id</i> <i>bpgrp-type-bpgrp-num</i> <i>bpgrp</i> keyword		

	type	ppp
	bpgrp-num	1 to 2000
aps-id	aps-group-id[.channel]	
	aps	keyword
	group-id	1 to 64
ccag-id	ccag-id.path-id [cc-type]	
	ccag	keyword
	id	1 to 8
	path-id	a, b
	cc-type	.sap-net, net-sap
eth-tunnel-id	eth-tunnel-id	
	eth-tunnel	keyword
	id	1 to 1024
lag-id	lag-id	
	lag	keyword
	id	1 to 800
gtg-id	gmpls-tun-grp-id	
	gmpls-tun-group	keyword
	id	1 to 1024

intermediate-destination-id — Displays information about the specified intermediate destination ID. 32 characters maximum.

no-inter-dest-id — Displays the information about no intermediate destination ID.

interface-name — Clears the interface name. 32 characters maximum.

dhcp

Syntax	dhcp
Context	clear>router
Description	This command enables the context to clear and reset DHCP entities.

statistics

Syntax	statistics [interface <i>ip-int-name</i> <i>ip-address</i>]
Context	clear>router>dhcp
Description	This command clears DHCP statistics.

forwarding-table

Syntax	forwarding-table [<i>slot-number</i>]				
Context	clear>router				
Description	This command clears the route table on the specified IOM with the route table. If the slot number is not specified, the command forces the route table to be recalculated.				
Parameters	<i>slot-number</i> — Clears the specified IOM slot. <table> <tr> <td>Values</td><td>1 to 10 (depending on chassis model)</td></tr> <tr> <td>Default</td><td>all IOMs</td></tr> </table>	Values	1 to 10 (depending on chassis model)	Default	all IOMs
Values	1 to 10 (depending on chassis model)				
Default	all IOMs				

interface

Syntax	interface [<i>ip-int-name</i> <i>ip-address</i>] [urpf-stats] [statistics] [hold-time] interface [<i>ip-int-name</i> <i>ip-address</i>] policy-accounting [class] [<i>index</i>] interface { <i>ip-int-name</i> <i>ip-address</i> } mac [<i>ieee-address</i>]								
Context	clear>router								
Description	This command clears IP interface statistics. If no IP interface is specified either by IP interface name or IP address, the command will perform the clear operation on all IP interfaces.								
Parameters	<i>ip-int-name</i> <i>ip-addr</i> — Specifies the IP interface name or IP interface address. <table> <tr> <td>Values</td><td>ip-int-name: 32 chars max ip-address: a.b.c.d</td></tr> <tr> <td>Default</td><td>all IP interfaces</td></tr> </table> urpf-stats — Resets the statistics associated with uRPF failures. statistics — Resets the IP interface traffic statistics. hold-time — Clears the IP interface activation hold time. policy-accounting — Clears the accounting statistics. class — Specifies whether to clear source class counters or destination class counters. <i>index</i> — Specifies the source-class or destination-class ID. <table> <tr> <td>Values</td><td>1 to 255</td></tr> </table> <i>ieee-address</i> — Specifies the MAC address. <table> <tr> <td>Values</td><td>xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx</td></tr> </table>	Values	ip-int-name: 32 chars max ip-address: a.b.c.d	Default	all IP interfaces	Values	1 to 255	Values	xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
Values	ip-int-name: 32 chars max ip-address: a.b.c.d								
Default	all IP interfaces								
Values	1 to 255								
Values	xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx								

bgp

Syntax	bgp
Context	clear>router>bgp
Description	This command clears or resets the route damping information for received routes.

damping

Syntax	damping [{ <i>ip-prefix</i> / <i>ip-prefix-length</i> [neighbor <i>ip-address</i>]}] [group <i>name</i>]
Context	clear>router>bgp
Description	This command clears or resets the route damping information for received routes.
Parameters	<i>ip-prefix/ip-prefix-length</i> — Clears damping information for entries that match the IP prefix and prefix length. Values ipv4-prefix: a.b.c.d (host bits must be 0) ipv4-prefix-le: 0 to 32 ipv6-prefix: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF] H • d: [0 to 255] D ipv6-prefix-le: 0 to 128 <i>ip-address</i> — Clears damping information for entries received from the BGP neighbor. Values ipv4-address: a.b.c.d 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 interface: 32 chars max, mandatory for link local addresses <i>name</i> — Clears damping information for entries received from any BGP neighbors in the peer group. 32 characters maximum.

flap-statistics

Syntax	flap-statistics [<i>ip-prefix</i> / <i>mask</i> [neighbor <i>ip-addr</i>]] group <i>group-name</i> regex <i>reg-exp</i> [policy <i>policy-name</i>]

Context	clear>router>bgp
Description	This command clears route flap statistics.
Parameters	<i>ip-prefix/mask</i> — Clears route flap statistics for entries that match the specified IP prefix and mask length. Values ipv4-prefix: a.b.c.d (host bits must be 0) ipv4-prefix-le: 0 to 32 ipv6-prefix: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF] H • d: [0 to 255] D ipv6-prefix-le: 0 to 128 <i>ip-addr</i> — Clears route flap statistics for entries received from the specified BGP neighbor. Values ipv4-prefix: a.b.c.d (host bits must be 0) 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 interface: 32 chars max, mandatory for link local addresses <i>group-name</i> — Clears route flap statistics for entries received from any BGP neighbors in the specified peer group. 32 characters maximum. <i>reg-exp</i> — Clears route flap statistics for all entries which have the regular expression and the AS path that matches the regular expression. 80 characters maximum. <i>policy-name</i> — Clears route flap statistics for entries that match the specified route policy. 32 characters maximum.

neighbor

Syntax	neighbor { <i>ip-address</i> as <i>as-number</i> external all } [soft soft-inbound statistics hard] neighbor <i>ip-address</i> as <i>as-number</i> external all } statistics neighbor <i>ip-address</i> end-of-rib
Context	clear>router>bgp
Description	This command resets the specified BGP peer or peers. This can cause existing BGP connections to be shutdown and restarted.

Parameters	<i>ip-address</i> — Resets the BGP neighbor with the specified IP address.
Values	ipv4-prefix: a.b.c.d (host bits must be 0) 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 interface: 32 chars max, mandatory for link local addresses
	<i>as-number</i> — Resets all BGP neighbors with the specified peer AS.
Values	1 to 4294967295
	external — Resets all EBGp neighbors.
	all — Resets all BGP neighbors.
	soft — The specified BGP neighbor(s) re-evaluates all routes in the Local-RIB against the configured export policies.
	soft-inbound — The specified BGP neighbor(s) re-evaluates all routes in the RIB-In against the configured import policies.
	statistics — The BGP neighbor statistics.
	hard — Resets the BGP session even when triggered-policy is used.
	end-of-rib — Clears the routing information base (RIB). This command applies when the router is helping the BGP neighbor through a BGP graceful restart. When the clear router bgp neighbor command is issued without the end-of-rib option and the neighbor is in the process of restarting, stale routes from the neighbor will be retained until the stale-routes-time is reached or else the neighbor exits graceful restart. When the command is issued with the end-of-rib option, stale routes from the neighbor are deleted immediately and graceful restart procedures are aborted.

protocol

Syntax	protocol
Context	clear>router>bgp
Description	This command resets the entire BGP protocol. If the AS number was previously changed, the BGP AS number does not inherit the new value.

database

Syntax	database
Context	clear>router>rip

Description This command flushes all routes in the RIP database.

statistics

Syntax **statistics** [**neighbor** {*ip-address* | *ip-int-name*}]

Context clear>router>rip

Description This command clears statistics for RIP neighbors.

Parameters {*ip-address* | *ip-int-name*} — Clears the statistics for the specified RIP interface.

Values ip-int-name: 32 characters maximum
ip-address: a.b.c.d

Default Clears statistics for all RIP interfaces.

id

Syntax **id** *service-id*

Context clear>service
clear>service>statistics

Description This command clears commands for a specific service.

Parameters *service-id* — The ID that uniquely identifies a service.

Values 1 to 2147483648

sap

Syntax **sap** *sap-id* {**all** | **counters** | **stp** | **l2pt** | **mrp**}

Context clear>service>statistics

Description This command clears SAP statistics for a SAP.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

all — Clears all SAP queue statistics and STP statistics.

counters — Clears all queue statistics associated with the SAP.

stp — Clears all STP statistics associated with the SAP.

l2pt — Clears all L2PT statistics associated with the SAP.

mrp — Clears all MRP statistics associated with the SAP.

dhcp

Syntax	dhcp
Context	clear>router>dhcp
Description	This command enables the context to clear DHCP parameters.

lease-state

Syntax	lease-state all [no-dhcp-release] lease-state [port port-id] inter-dest-id intermediate-destination-id [no-dhcp-release] lease-state [port port-id] no-inter-dest-id [no-dhcp-release] lease-state ip-address ip-address [/mask] [no-dhcp-release] lease-state mac ieee-address [no-dhcp-release] lease-state port port-id [no-dhcp-release] lease-state sap sap-id [no-dhcp-release] lease-state sdp sdp-id:vc-id [no-dhcp-release]
Context	clear>service>id>dhcp
Description	This command clears DHCP lease state information for this service.
Parameters	all — Clears all lease state statistics. <i>port-id</i> — Displays information about the specified port ID. <i>intermediate-destination-id</i> — Displays information about the specified intermediate destination ID. 32 characters maximum. no-inter-dest-id — Displays the information about no intermediate destination ID. <i>ip-address[/mask]</i> — The IP address of the IP interface. The <i>ip-address</i> portion of the address command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets). Values a.b.c.d. mask: 1 to 32 <i>ieee-address</i> — 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. Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition.

sdp-id — The SDP ID to be cleared.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID to be cleared.

Values 1 to 4294967295

no-dhcp-release — Specifies that the node will clear the state without sending the DHCP release message.

site

Syntax **site** *service-id*

Context clear>service>id

Description This command clears site-specific information for the service.

Parameters *service-id* — Specifies the service ID or service name up to 64 characters long.
Values 1 to 2147483648

spoke-sdp

Syntax **spoke-sdp** *sdp-id:vc-id* **ingress-vc-label**
spoke-sdp *sdp-id:vc-id* **l2tpv3**
spoke-sdp *sdp-id:vc-id* **vccv-bfd** {**session** | **statistics**}

Context clear>service>id

Description This command clears and resets the spoke SDP bindings for the service.

Parameters *sdp-id* — The spoke SDP ID to be reset.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID to be reset.

Values 1 to 4294967295

ingress-vc-label — Specifies to clear the ingress VC label.

l2tpv3 — Specifies to clear the session mismatch flag on the spoke-SDP binding after the flag was set to true by a detected mismatch between the configured parameters and the received parameters.

vccv-bfd — Specifies to clear the following information related to VCCV BFD on the spoke-sdp.

session — Clears session information.

statistics — Clears spoke SDP statistics.

sdp

Syntax	sdp <i>sdp-id</i> keep-alive
Context	clear>service>statistics
Description	This command clears keepalive statistics associated with the SDP ID.
Parameters	<i>sdp-id</i> — The SDP ID for which to clear keepalive statistics. Values 1 to 17407 keep-alive — Clears the keep-alive history associated with this SDP ID.

counters

Syntax	counters
Context	clear>service>statistics>id
Description	Clears all traffic queue counters associated with the service ID.

spoke-sdp

Syntax	spoke-sdp <i>sdp-id[:vc-id]</i> {all counters stp l2pt mrp}
Context	clear>service>statistics>id
Description	This command clears statistics for the spoke SDP bound to the service.
Parameters	<i>sdp-id</i> — The spoke SDP ID for which to clear statistics. Values 1 to 17407 <i>vc-id</i> — The virtual circuit ID on the SDP ID to be reset. Values 1 to 4294967295 all — Clears all queue statistics and STP statistics associated with the SDP. counters — Clears all queue statistics associated with the SDP. stp — Clears all STP statistics associated with the SDP. l2pt — Clears all L2PT statistics associated with the SDP. mrp — Clears all MRP statistics associated with the SDP.

stp

Syntax	stp

Context	clear>service>statistics>id
Description	This command clears all spanning tree statistics for the service ID.

3.9.2.3 VPRN Debug Commands

id

Syntax	[no] id <i>service-id</i>
Context	debug>service
Description	This command debugs commands for a specific service. The no form of the command disables debugging.
Parameters	<i>service-id</i> — The ID that uniquely identifies a service.

arp-host

Syntax	[no] arp-host
Context	debug>service>id
Description	This command enables and configures ARP host debugging. The no form of the command disables ARP host debugging.

dhcp

Syntax	[no] dhcp
Context	debug>service>id
Description	This command enables the context for DHCP debugging. The no form of the command disables DHCP debugging.

detail-level

Syntax	detail-level {low medium high} no detail-level

Context	debug>service>id>dhcp
Description	This command configures the DHCP tracing detail level. The no form of the command disables debugging.
Parameters	low — Displays a low detail level for DHCP debugging. medium — Displays a medium detail level for DHCP debugging. high — Displays a high detail level for DHCP debugging.

mode

Syntax	mode {dropped-only ingr-and-dropped egr-ingr-and-dropped} no mode
Context	debug>service>id>dhcp
Description	This command configures the DHCP tracing mode. The no form of the command disables debugging.
Parameters	dropped-only — Only displays dropped packets. ingr-and-dropped — Only displays ingress packet and dropped packets. egr-ingr-and-dropped — Displays ingress, egress and dropped packets.

host-connectivity-verify

Syntax	[no] host-connectivity-verify
Context	debug>service>id
Description	This command enables Subscriber Host Connectivity Verification (SHCV) debugging. The no form of the command disables the SHCV debugging.

ip

Syntax	[no] ip <i>ip-address</i>
Context	debug>service>id>host-connectivity-verify
Description	This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular IP address.

Parameters	<i>ip-address</i> — The IP address of the IP interface. The <i>ip-address</i> portion of the address command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets). Values ipv4-prefix: a.b.c.d (host bits must be 0) 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

mac

Syntax	[no] mac <i>ieee-address</i>
Context	debug>service>id>host-connectivity-verify
Description	This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular MAC address.
Parameters	<i>mac-address</i> — 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. Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx (cannot be all zeros)

sap

Syntax	[no] sap <i>sap-id</i>
Context	debug>service>id>host-connectivity-verify
Description	This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular SAP.
Parameters	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition.

sap

Syntax	[no] sap <i>sap-id</i>
Context	debug>service>id debug>service>id>dhcp

debug>service>stp

Description This command enables STP debugging for a specific SAP.

The **no** form of the command disables debugging.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

sdp

Syntax [no] **sdp** *sdp-id:vc-id*

Context debug>service>id
debug>service>id>dhcp
debug>service>stp

Description This command enables STP debugging for a specific SDP.

The **no** form of the command disables debugging.

Parameters *sdp-id:vc-id* — Specifies the SDP ID and VC ID.

Values sdp-id: 1 to 17407
vc-id: 1 to 4294967295

event-type

Syntax [no] **event-type** {**config-change** | **svc-oper-status-change** | **sap-oper-status-change** | **sdpbind-oper-status-change**}

Context debug>service>id

Description This command enables debugging for a particular event type.

The **no** form of the command disables debugging.

event-type

Syntax [no] **event-type** {**config-change** | **oper-status-change** | **neighbor discovery** | **control-channel-status**}

Context debug>service>id>sdp

Description This command enables debugging for a particular event type.

The **no** form of the command disables debugging.

event-type

Syntax	[no] event-type {arp config-change oper-status-change neighbor-discovery}
Context	debug>service>id>sap
Description	This command enables debugging for a particular event type. The no form of the command disables debugging.

stp

Syntax	[no] stp
Context	debug>service>id
Description	This command enables the context for debugging STP. The no form of the command disables debugging.

all-events

Syntax	all-events
Context	debug>service>id>event-type
Description	This command enables STP debugging for all events. The no form of the command disables debugging.

bpdu

Syntax	[no] bpdu
Context	debug>service>stp
Description	This command enables STP debugging for received and transmitted BPDUs. The no form of the command disables debugging.

core-connectivity

Syntax	[no] core-connectivity
Context	debug>service>stp

Description This command enables STP debugging for core connectivity.
The **no** form of the command disables debugging.

exception

Syntax [no] exception
Context debug>service>stp
Description This command enables STP debugging for exceptions.
The **no** form of the command disables debugging.

fsm-state-changes

Syntax [no] fsm-state-changes
Context debug>service>stp
Description This command enables STP debugging for FSM state changes.
The **no** form of the command disables debugging.

fsm-timers

Syntax [no] fsm-timers
Context debug>service>stp
Description This command enables STP debugging for FSM timer changes.
The **no** form of the command disables debugging.

port-role

Syntax [no] port-role
Context debug>service>stp
Description This command enables STP debugging for changes in port roles.
The **no** form of the command disables debugging.

port-state


Syntax	[no] port-state
Context	debug>service>stp
Description	This command enables STP debugging for port states. The no form of the command disables debugging.

3.10 Tools Command Reference

3.10.1 Command Hierarchies

- [Tools Commands](#)

3.10.1.1 Tools Commands



3.10.2 Command Descriptions

3.10.2.1 Tools Commands

tools

Syntax	tools
Context	root
Description	This command enables the context to configure or display tools.

dump

Syntax	dump
Context	tools

Description	This command enables the context to display tools information.

perform

Syntax	perform
Context	tools
Description	This command enables the context to configure tools that perform specific tasks.

service

Syntax	service
Context	tools>perform tools>dump
Description	This command enables the context to configure and view tools for service debugging purposes.

id

Syntax	id <i>service-id</i>
Context	tools>dump>service tools>perform>service
Description	This command specifies a service for which to enable and view tools for service debugging purposes.
Parameters	<i>service-id</i> — Specifies the service ID.

interface

Syntax	interface <i>ip-int-name</i>
Context	tools>dump>service>id tools>perform>service>id
Description	This command specifies an IP interface for which to enable and view tools for service debugging purposes.
Parameters	<i>ip-int-name</i> — The name of the IP interface. 32 characters maximum.

ignore-sap-port-state

Syntax	ignore-sap-port-state
Context	tools>dump>service>id>interface
Description	This command displays all service interfaces that have accepted an ignore-sap-port-state start command.

This command can be executed without a *service-id* for a complete list of interfaces that have accepted an **ignore-sap-port-state start** command. The command can be executed within a specific *service-id* context for a list of all interfaces for the specified service that have accepted the **ignore-sap-port-state start** command. The *ip-int-name* parameter may be optionally configured to display results only for the specified interface. If the command is executed against a specific interface that has not accepted an **ignore-sap-port-state start** command, the display command will display a message indicating that no action has been started for the interface.

Output The following output is an example of **ignore-sap-port-state** information.

Sample Output

```
tools dump service ignore-sap-port-state
=====
Interface Marked Ignore SAP Operational State
=====
SvcId      Interface-Name      Adm/Opr (v4/v6)  Type  Port/SapId
          IP-Address                               PfxState
-----
1001      ies-1001            Up/ (Up/Down)    IES    1/1/1:1001.
          192.168.3.30/24    n/a
-----
Number of entries: 1
=====
```

ignore-sap-port-state

Syntax	ignore-sap-port-state {start stop}
Context	tools>perform>service>id>interface
Description	This command enables bypassing of the Ethernet operational state check, which would otherwise be part of the Ethernet SAP operational state checking function. All other checks are performed as normal.

This command may be executed against IES and VPRN service IP interfaces directly configured over an Ethernet SAP. When the command is executed against an operational Ethernet SAP, the command enters a pending state, waiting for a non-operational change. Network interfaces have no SAP association and do not support this feature. When using **subscriber-interface** and **group-interface**, the command is only applicable to the **group-interface** associated with the SAP. R-VPLS does not have Ethernet SAPs directly configured under the interface, and is not supported.

This function is meant to allow service validation and reachability testing when a physical Ethernet port has not been connected. The command may be executed for a non-operational SAP that is cabled. However, if the SAP transitions to an operational state, ingress and egress packet processing may still occur.

This command configuration does not survive a system restart.

- Parameters**
- start** — Enables port state bypass mode for the interface. If the Ethernet port SAP is already operational, there is no immediate effect of the command, and the Ignore Port state under the **show service id all** command will show “pending”. However, if the interface Ethernet port transitions to a non-operational state, the command then bypasses this port state and executes the remainder of the operational checks, and the flag for the Ignore Port state transitions to “active”. When this command is in effect, the SAP represents the state of the SAP ignoring the port state.
 - stop** — Disables port state bypass mode for the interface attached to the SAP.

4 Standards and Protocol Support



Note: The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

Access Node Control Protocol (ANCP)

draft-ietf-ancp-protocol-02, *Protocol for Access Node Control Mechanism in Broadband Networks*

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

Application Assurance (AA)

3GPP Release 12 (ADC rules over Gx interfaces)

RFC 3507, *Internet Content Adaptation Protocol (ICAP)*

Asynchronous Transfer Mode (ATM)

AF-ILMI-0065.000, *Integrated Local Management Interface (ILMI) Version 4.0*

AF-PHY-0086.001, *Inverse Multiplexing for ATM (IMA) Specification Version 1.1*

AF-TM-0121.000, *Traffic Management Specification Version 4.1*

AF-TM-0150.00, *Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR*

GR-1113-CORE, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1*

GR-1248-CORE, *Generic Requirements for Operations of ATM Network Elements (NEs), Issue 3*

ITU-T I.432.1, *B-ISDN user-network interface - Physical layer specification: General characteristics (02/99)*

ITU-T I.610, *B-ISDN operation and maintenance principles and functions (11/95)*

RFC 1626, *Default IP MTU for use over ATM AAL5*

RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*

Bidirectional Forwarding Detection (BFD)

RFC 5880, *Bidirectional Forwarding Detection (BFD)*

RFC 5881, *Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop)*

RFC 5883, *Bidirectional Forwarding Detection (BFD) for Multihop Paths*

RFC 7130, *Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces*

Border Gateway Protocol (BGP)

draft-hares-idr-update-attr-low-bits-fix-01, *Update Attribute Flag Low Bits Clarification*

draft-ietf-idr-add-paths-guidelines-08, *Best Practices for Advertisement of Multiple Paths in IBGP*

draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*

draft-ietf-idr-bgp-flowspec-oid-03, *Revised Validation Procedure for BGP Flow Specifications*

draft-ietf-idr-bgp-gr-notification-01, *Notification Message support for BGP Graceful Restart*

draft-ietf-idr-bgp-optimal-route-reflection-10, *BGP Optimal Route Reflection (BGP-ORR)*

draft-ietf-idr-error-handling-03, *Revised Error Handling for BGP UPDATE Messages*

draft-ietf-idr-flowspec-interfaceset-03, *Applying BGP flowspec rules on a specific interface set*

draft-ietf-idr-flowspec-path-redirect-05, *Flowspec Indirection-id Redirect (localised ID)*

draft-ietf-idr-flowspec-redirect-ip-02, *BGP Flow-Spec Redirect to IP Action*

draft-ietf-idr-link-bandwidth-03, *BGP Link Bandwidth Extended Community*

draft-ietf-sidr-origin-validation-signaling-04, *BGP Prefix Origin Validation State Extended Community*

draft-uttaro-idr-bgp-persistence-03, *Support for Long-lived BGP Graceful Restart*

RFC 1772, *Application of the Border Gateway Protocol in the Internet*

RFC 1997, *BGP Communities Attribute*

RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*

RFC 2439, *BGP Route Flap Damping*

RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*

RFC 2858, *Multiprotocol Extensions for BGP-4*

RFC 2918, *Route Refresh Capability for BGP-4*

RFC 3107, *Carrying Label Information in BGP-4*

RFC 3392, *Capabilities Advertisement with BGP-4*

RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*

RFC 4360, *BGP Extended Communities Attribute*

RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 4456, *BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)*

RFC 4486, *Subcodes for BGP Cease Notification Message*
RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*
RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/
MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual
Private Networks (VPNs)*
RFC 4724, *Graceful Restart Mechanism for BGP (helper mode)*
RFC 4760, *Multiprotocol Extensions for BGP-4*
RFC 4798, *Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge
Routers (6PE)*
RFC 4893, *BGP Support for Four-octet AS Number Space*
RFC 5004, *Avoid BGP Best Path Transitions from One External to Another*
RFC 5065, *Autonomous System Confederations for BGP*
RFC 5291, *Outbound Route Filtering Capability for BGP-4*
RFC 5396, *Textual Representation of Autonomous System (AS) Numbers (asplain)*
RFC 5549, *Advertising IPv4 Network Layer Reachability Information with an IPv6
Next Hop*
RFC 5575, *Dissemination of Flow Specification Rules*
RFC 5668, *4-Octet AS Specific BGP Extended Community*
RFC 6286, *Autonomous-System-Wide Unique BGP Identifier for BGP-4*
RFC 6810, *The Resource Public Key Infrastructure (RPKI) to Router Protocol*
RFC 6811, *Prefix Origin Validation*
RFC 6996, *Autonomous System (AS) Reservation for Private Use*
RFC 7311, *The Accumulated IGP Metric Attribute for BGP*
RFC 7607, *Codification of AS 0 Processing*
RFC 7674, *Clarification of the Flowspec Redirect Extended Community*
RFC 7752, *North-Bound Distribution of Link-State and Traffic Engineering (TE)
Information Using BGP*
RFC 7854, *BGP Monitoring Protocol (BMP)*
RFC 7911, *Advertisement of Multiple Paths in BGP*
RFC 7999, *BLACKHOLE Community*
RFC 8092, *BGP Large Communities Attribute*

Circuit Emulation

RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet
(SAToP)*
RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation
Service over Packet Switched Network (CESoPSN)*

RFC 5287, *Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks*

Ethernet

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*

IEEE 802.1ad, *Provider Bridges*

IEEE 802.1ag, *Connectivity Fault Management*

IEEE 802.1ah, *Provider Backbone Bridges*

IEEE 802.1ak, *Multiple Registration Protocol*

IEEE 802.1aq, *Shortest Path Bridging*

IEEE 802.1ax, *Link Aggregation*

IEEE 802.1D, *MAC Bridges*

IEEE 802.1p, *Traffic Class Expediting*

IEEE 802.1Q, *Virtual LANs*

IEEE 802.1s, *Multiple Spanning Trees*

IEEE 802.1w, *Rapid Reconfiguration of Spanning Tree*

IEEE 802.1X, *Port Based Network Access Control*

IEEE 802.3ab, *1000BASE-T*

IEEE 802.3ac, *VLAN Tag*

IEEE 802.3ad, *Link Aggregation*

IEEE 802.3ae, *10 Gb/s Ethernet*

IEEE 802.3ah, *Ethernet in the First Mile*

IEEE 802.3ba, *40 Gb/s and 100 Gb/s Ethernet*

IEEE 802.3i, *Ethernet*

IEEE 802.3u, *Fast Ethernet*

IEEE 802.3x, *Ethernet Flow Control*

IEEE 802.3z, *Gigabit Ethernet*

ITU-T G.8031/Y.1342, *Ethernet Linear Protection Switching*

ITU-T G.8032/Y.1344, *Ethernet Ring Protection Switching*

ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

Ethernet VPN (EVPN)

draft-ietf-bess-evpn-ac-df-01, *AC-Influenced Designated Forwarder Election for EVPN*

draft-ietf-bess-evpn-pref-df-01, *Preference-based EVPN DF Election*

draft-ietf-bess-evpn-prefix-advertisement-11, *IP Prefix Advertisement in EVPN*

draft-ietf-bess-evpn-proxy-arp-nd-04, *Operational Aspects of Proxy-ARP/ND in EVPN Networks*
draft-ietf-bess-evpn-vpls-seamless-integ-03, *(PBB-)EVPN Seamless Integration with (PBB-)VPLS*
draft-snr-bess-pbb-evpn-isid-cmacflush-01, *PBB-EVPN ISID-based CMAC-Flush*
RFC 7432, *BGP MPLS-Based Ethernet VPN*
RFC 7623, *Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)*
RFC 8214, *Virtual Private Wire Service Support in Ethernet VPN*
RFC 8317, *Ethernet-Tree (E-Tree) Support in Ethernet VPN (EVPN) an Provider Backbone Bridging EVPN (PBB-EVPN)*
RFC 8365, *A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)*

Frame Relay

ANSI T1.617 Annex D, *DSS1 - Signalling Specification For Frame Relay Bearer Service*
FRF.1.2, *PVC User-to-Network Interface (UNI) Implementation Agreement*
FRF.12, *Frame Relay Fragmentation Implementation Agreement*
FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*
FRF.5, *Frame Relay/ATM PVC Network Interworking Implementation*
FRF2.2, *PVC Network-to-Network Interface (NNI) Implementation Agreement*
ITU-T Q.933 Annex A, *Additional procedures for Permanent Virtual Connection (PVC) status management*

Generalized Multiprotocol Label Switching (GMPLS)

draft-ietf-ccamp-rsvp-te-srlg-collect-04, *RSVP-TE Extensions for Collecting SRLG Information*
RFC 3471, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description*
RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions*
RFC 4204, *Link Management Protocol (LMP)*
RFC 4208, *Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model*
RFC 4872, *RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery*
RFC 5063, *Extensions to GMPLS Resource Reservation Protocol (RSVP) Graceful Restart (helper mode)*

gRPC Remote Procedure Calls (gRPC)

gnmi.proto, *gRPC Network Management Interface (gNMI), version 0.4.0*

*gRPC Network Management Interface (gNMI), Capabilities, Get, Set, Subscribe
(ONCE, SAMPLE, ON_CHANGE)*

Intermediate System to Intermediate System (IS-IS)

draft-ietf-isis-mi-02, *IS-IS Multi-Instance*

draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*

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)*

RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*

RFC 2973, *IS-IS Mesh Groups*

RFC 3359, *Reserved Type, Length and Value (TLV) Codepoints in Intermediate
System to Intermediate System*

RFC 3719, *Recommendations for Interoperable Networks using Intermediate
System to Intermediate System (IS-IS)*

RFC 3787, *Recommendations for Interoperable IP Networks using Intermediate
System to Intermediate System (IS-IS)*

RFC 4971, *Intermediate System to Intermediate System (IS-IS) Extensions for
Advertising Router Information*

RFC 5120, *M-ISIS: Multi Topology (MT) Routing in IS-IS*

RFC 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*

RFC 5301, *Dynamic Hostname Exchange Mechanism for IS-IS*

RFC 5302, *Domain-wide Prefix Distribution with Two-Level IS-IS*

RFC 5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*

RFC 5304, *IS-IS Cryptographic Authentication*

RFC 5305, *IS-IS Extensions for Traffic Engineering TE*

RFC 5306, *Restart Signaling for IS-IS (helper mode)*

RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label
Switching (GMPLS)*

RFC 5308, *Routing IPv6 with IS-IS*

RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*

RFC 5310, *IS-IS Generic Cryptographic Authentication*

RFC 6213, *IS-IS BFD-Enabled TLV*

RFC 6232, *Purge Originator Identification TLV for IS-IS*

RFC 6233, *IS-IS Registry Extension for Purges*

RFC 6329, *IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging*
RFC 7775, *IS-IS Route Preference for Extended IP and IPv6 Reachability*
RFC 7794, *IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability*
RFC 8202, *IS-IS Multi-Instance* (single topology)

Internet Protocol (IP) — Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*
RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*
RFC 7431, *Multicast-Only Fast Reroute*
RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*

Internet Protocol (IP) — General

draft-grant-tacacs-02, *The TACACS+ Protocol*
RFC 768, *User Datagram Protocol*
RFC 793, *Transmission Control Protocol*
RFC 854, *Telnet Protocol Specifications*
RFC 1350, *The TFTP Protocol (revision 2)*
RFC 2347, *TFTP Option Extension*
RFC 2348, *TFTP Blocksize Option*
RFC 2349, *TFTP Timeout Interval and Transfer Size Options*
RFC 2428, *FTP Extensions for IPv6 and NATs*
RFC 2784, *Generic Routing Encapsulation (GRE)*
RFC 2890, *Key and Sequence Number Extensions to GRE*
RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*
RFC 4251, *The Secure Shell (SSH) Protocol Architecture*
RFC 4252, *The Secure Shell (SSH) Authentication Protocol* (publickey, password)
RFC 4253, *The Secure Shell (SSH) Transport Layer Protocol*
RFC 4254, *The Secure Shell (SSH) Connection Protocol*
RFC 4632, *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*
RFC 5082, *The Generalized TTL Security Mechanism (GTSM)*
RFC 5656, *Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer (ECDSA)*
RFC 5925, *The TCP Authentication Option*
RFC 5926, *Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)*
RFC 6398, *IP Router Alert Considerations and Usage (MLD)*

RFC 6528, *Defending against Sequence Number Attacks*

Internet Protocol (IP) — Multicast

cisco-ipmulticast/pim-autorp-spec01, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast* (version 1)

draft-dolganow-bess-mvpn-expl-track-01, *Explicit Tracking with Wild Card Routes in Multicast VPN*

draft-ietf-bier-mvpn-11, *Multicast VPN Using BIER*

draft-ietf-idmr-traceroute-ipm-07, *A "traceroute" facility for IP Multicast*

draft-ietf-l2vpn-vpls-pim-snooping-07, *Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)*

draft-ietf-mboned-mtrace-v2-17, *Mtrace Version 2: Traceroute Facility for IP Multicast*

RFC 1112, *Host Extensions for IP Multicasting*

RFC 2236, *Internet Group Management Protocol, Version 2*

RFC 2365, *Administratively Scoped IP Multicast*

RFC 2375, *IPv6 Multicast Address Assignments*

RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*

RFC 3306, *Unicast-Prefix-based IPv6 Multicast Addresses*

RFC 3376, *Internet Group Management Protocol, Version 3*

RFC 3446, *Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*

RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*

RFC 3618, *Multicast Source Discovery Protocol (MSDP)*

RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*

RFC 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*

RFC 3973, *Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)* (auto-RP groups)

RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*

RFC 4604, *Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast*

RFC 4607, *Source-Specific Multicast for IP*

RFC 4608, *Source-Specific Protocol Independent Multicast in 232/8*

RFC 4610, *Anycast-RP Using Protocol Independent Multicast (PIM)*

RFC 4611, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios*

RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*

RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3) / Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction*

RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*

RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*

RFC 6037, *Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs*

RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*

RFC 6513, *Multicast in MPLS/BGP IP VPNs*

RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/IP VPNs*

RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs*

RFC 6516, *IPv6 Multicast VPN (MVPN) Support Using PIM Control Plane and Selective Provider Multicast Service Interface (S-PMSI) Join Messages*

RFC 6625, *Wildcards in Multicast VPN Auto-Discover Routes*

RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path*

RFC 7246, *Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context*

RFC 7385, *IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points*

RFC 7716, *Global Table Multicast with BGP Multicast VPN (BGP-MVPN) Procedures*

RFC 7761, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*

RFC 8279, *Multicast Using Bit Index Explicit Replication (BIER)*

RFC 8296, *Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks (MPLS encapsulation)*

RFC 8401, *Bit Index Explicit Replication (BIER) Support via IS-IS*

Internet Protocol (IP) — Version 4

RFC 791, *Internet Protocol*

RFC 792, *Internet Control Message Protocol*

RFC 826, *An Ethernet Address Resolution Protocol*

RFC 951, *Bootstrap Protocol (BOOTP)*

RFC 1034, *Domain Names - Concepts and Facilities*

RFC 1035, *Domain Names - Implementation and Specification*

RFC 1191, *Path MTU Discovery (router specification)*

RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*
RFC 1534, *Interoperation between DHCP and BOOTP*
RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*
RFC 1812, *Requirements for IPv4 Routers*
RFC 1918, *Address Allocation for Private Internets*
RFC 2003, *IP Encapsulation within IP*
RFC 2131, *Dynamic Host Configuration Protocol*
RFC 2132, *DHCP Options and BOOTP Vendor Extensions*
RFC 2401, *Security Architecture for Internet Protocol*
RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*
RFC 3046, *DHCP Relay Agent Information Option (Option 82)*
RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*
RFC 4884, *Extended ICMP to Support Multi-Part Messages (ICMPv4 and ICMPv6 Time Exceeded)*

Internet Protocol (IP) — Version 6

RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*
RFC 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*
RFC 3122, *Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification*
RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*
RFC 3587, *IPv6 Global Unicast Address Format*
RFC 3596, *DNS Extensions to Support IP version 6*
RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*
RFC 3646, *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*
RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*
RFC 3971, *SEcure Neighbor Discovery (SEND)*
RFC 3972, *Cryptographically Generated Addresses (CGA)*
RFC 4007, *IPv6 Scoped Address Architecture*
RFC 4193, *Unique Local IPv6 Unicast Addresses*
RFC 4291, *Internet Protocol Version 6 (IPv6) Addressing Architecture*
RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*
RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*
RFC 4862, *IPv6 Stateless Address Autoconfiguration (router functions)*

RFC 4890, *Recommendations for Filtering ICMPv6 Messages in Firewalls*
RFC 4941, *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*
RFC 5007, *DHCPv6 Leasequery*
RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*
RFC 5722, *Handling of Overlapping IPv6 Fragments*
RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6 (IPv6)*
RFC 5952, *A Recommendation for IPv6 Address Text Representation*
RFC 6092, *Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service* (Internet Control and Management, Upper-Layer Transport Protocols, UDP Filters, IPsec and Internet Key Exchange (IKE), TCP Filters)
RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*
RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*
RFC 8021, *Generation of IPv6 Atomic Fragments Considered Harmful*
RFC 8200, *Internet Protocol, Version 6 (IPv6) Specification*
RFC 8201, *Path MTU Discovery for IP version 6*

Internet Protocol Security (IPsec)

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*
draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*
RFC 2401, *Security Architecture for the Internet Protocol*
RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*
RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*
RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*
RFC 2406, *IP Encapsulating Security Payload (ESP)*
RFC 2407, *IPsec Domain of Interpretation for ISAKMP (IPsec DoI)*
RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)*
RFC 2409, *The Internet Key Exchange (IKE)*
RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*
RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman group for Internet Key Exchange (IKE)*
RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*
RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*
RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*
RFC 3947, *Negotiation of NAT-Traversal in the IKE*

RFC 3948, *UDP Encapsulation of IPsec ESP Packets*
RFC 4106, *The Use of Galois/Counter Mode (GCM) in IPsec ESP*
RFC 4210, *Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)*
RFC 4211, *Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)*
RFC 4301, *Security Architecture for the Internet Protocol*
RFC 4303, *IP Encapsulating Security Payload*
RFC 4307, *Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*
RFC 4308, *Cryptographic Suites for IPsec*
RFC 4434, *The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)*
RFC 4543, *The Use of Galois Message Authentication Code (GMAC) in IPsec ESP and AH*
RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPSec*
RFC 4945, *The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2 and PKIX*
RFC 5019, *The Lightweight Online Certificate Status Protocol (OCSP) Profile for High-Volume Environments*
RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*
RFC 5282, *Using Authenticated Encryption Algorithms with the Encrypted Payload of the IKEv2 Protocol*
RFC 5903, *ECP Groups for IKE and IKEv2*
RFC 5998, *An Extension for EAP-Only Authentication in IKEv2*
RFC 6379, *Suite B Cryptographic Suites for IPsec*
RFC 6380, *Suite B Profile for Internet Protocol Security (IPsec)*
RFC 6712, *Internet X.509 Public Key Infrastructure -- HTTP Transfer for the Certificate Management Protocol (CMP)*
RFC 6960, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*
RFC 7296, *Internet Key Exchange Protocol Version 2 (IKEv2)*
RFC 7321, *Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)*
RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*
RFC 7427, *Signature Authentication in the Internet Key Exchange Version 2 (IKEv2)*
RFC 7468, *Textual Encodings of PKIX, PKCS, and CMS Structures*

Label Distribution Protocol (LDP)

draft-ietf-mpls-ldp-ip-pw-capability-09, Controlling State Advertisements Of Non-negotiated LDP Applications
draft-pdutta-mpls-ldp-adj-capability-00, LDP Adjacency Capabilities
draft-pdutta-mpls-ldp-v2-00, LDP Version 2
draft-pdutta-mpls-mldp-up-redundancy-00, Upstream LSR Redundancy for Multipoint LDP Tunnels
draft-pdutta-mpls-multi-ldp-instance-00, Multiple LDP Instances
draft-pdutta-mpls-tldp-hello-reduce-04, Targeted LDP Hello Reduction
RFC 3037, LDP Applicability
RFC 3478, Graceful Restart Mechanism for Label Distribution Protocol (helper mode)
RFC 5036, LDP Specification
RFC 5283, LDP Extension for Inter-Area Label Switched Paths (LSPs)
RFC 5443, LDP IGP Synchronization
RFC 5561, LDP Capabilities
RFC 5919, Signaling LDP Label Advertisement Completion
RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths
RFC 6512, Using Multipoint LDP When the Backbone Has No Route to the Root
RFC 6826, Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths
RFC 7032, LDP Downstream-on-Demand in Seamless MPLS
RFC 7552, Updates to LDP for IPv6

Layer Two Tunneling Protocol (L2TP) Network Server (LNS)

draft-mammoliti-l2tp-accessline-avp-04, Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions
RFC 2661, Layer Two Tunneling Protocol "L2TP"
RFC 2809, Implementation of L2TP Compulsory Tunneling via RADIUS
RFC 3438, Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers: Internet Assigned Numbers Authority (IANA) Considerations Update
RFC 3931, Layer Two Tunneling Protocol - Version 3 (L2TPv3)
RFC 4719, Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)
RFC 4951, Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"

Management

draft-ietf-snmpv3-update-mib-05, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*

draft-ietf-isis-wg-mib-06, *Management Information Base for Intermediate System to Intermediate System (IS-IS)*

draft-ietf-mboned-msdp-mib-01, *Multicast Source Discovery protocol MIB*

draft-ietf-mpls-ldp-mib-07, *Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP)*

draft-ietf-mpls-lsr-mib-06, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2*

draft-ietf-mpls-te-mib-04, *Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base*

draft-ietf-ospf-mib-update-08, *OSPF Version 2 Management Information Base*

draft-ietf-vrrp-unified-mib-06, *Definitions of Managed Objects for the VRRP over IPv4 and IPv6 (IPv6)*

ianaaddressfamilynumbers-mib, *IANA-ADDRESS-FAMILY-NUMBERS-MIB*

ianagmplstc-mib, *IANA-GMPLS-TC-MIB*

ianaiftype-mib, *IANAifType-MIB*

ianaiprouteprotocol-mib, *IANA-RTPROTO-MIB*

IEEE8021-CFM-MIB, *IEEE P802.1ag(TM) CFM MIB*

IEEE8021-PAE-MIB, *IEEE 802.1X MIB*

IEEE8023-LAG-MIB, *IEEE 802.3ad MIB*

LLDP-MIB, *IEEE P802.1AB(TM) LLDP MIB*

openconfig-bgp.yang version 3.0.1, *BGP Module*

openconfig-bgp-common.yang version 3.0.1, *BGP Common Module*

openconfig-bgp-common-multiprotocol.yang version 3.0.1, *BGP Common Multiprotocol Module*

openconfig-bgp-common-structure.yang version 3.0.1, *BGP Common Structure Module*

openconfig-bgp-global.yang version 3.0.1, *BGP Global Module*

openconfig-bgp-neighbor.yang version 3.0.1, *BGP Neighbor Module*

openconfig-bgp-peer-group.yang version 3.0.1, *BGP Peer Group Module*

openconfig-bgp-policy.yang version 4.0.1, *BGP Policy Module*

openconfig-if-aggregate.yang version 2.0.0, *Interfaces Aggregated Model*

openconfig-if-ethernet.yang version 2.0.0, *Interfaces Ethernet Model*

openconfig-if-ip.yang version 2.0.0, *Interfaces IP Module*

openconfig-if-ip-ext.yang version 2.0.0, *Interfaces IP Extensions Module*

openconfig-interfaces.yang version 2.0.0, *Interfaces Module*

openconfig-isis.yang version 0.3.0, *IS-IS Module*
openconfig-isis-lsp.yang version 0.3.0, *IS-IS LSP Module*
openconfig-isis-routing.yang version 0.3.0, *IS-IS Routing Module*
openconfig-lacp.yang version 1.1.0, *LACP Module*
openconfig-lldp.yang version 0.1.0, *LLDP Module*
openconfig-local-routing.yang version 1.0.1, *Local Routing Module*
openconfig-network-instance.yang version 0.8.0, *Network Instance Module*
openconfig-routing-policy.yang version 3.0.0, *Routing Policy Module*
openconfig-vlan.yang version 2.0.0, *VLAN Module*
RFC 1157, *A Simple Network Management Protocol (SNMP)*
RFC 1212, *Concise MIB Definitions*
RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*
RFC 1215, *A Convention for Defining Traps for use with the SNMP*
RFC 1724, *RIP Version 2 MIB Extension*
RFC 1901, *Introduction to Community-based SNMPv2*
RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIv2*
RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIv2*
RFC 2206, *RSVP Management Information Base using SMIv2*
RFC 2213, *Integrated Services Management Information Base using SMIv2*
RFC 2494, *Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type*
RFC 2514, *Definitions of Textual Conventions and OBJECT-IDENTITIES for ATM Management*
RFC 2515, *Definitions of Managed Objects for ATM Management*
RFC 2570, *SNMP Version 3 Framework*
RFC 2571, *An Architecture for Describing SNMP Management Frameworks*
RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*
RFC 2573, *SNMP Applications*
RFC 2574, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*
RFC 2575, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*
RFC 2578, *Structure of Management Information Version 2 (SMIv2)*
RFC 2579, *Textual Conventions for SMIv2*
RFC 2580, *Conformance Statements for SMIv2*

RFC 2787, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol*

RFC 2819, *Remote Network Monitoring Management Information Base*

RFC 2856, *Textual Conventions for Additional High Capacity Data Types*

RFC 2863, *The Interfaces Group MIB*

RFC 2864, *The Inverted Stack Table Extension to the Interfaces Group MIB*

RFC 2933, *Internet Group Management Protocol MIB*

RFC 3014, *Notification Log MIB*

RFC 3164, *The BSD syslog Protocol*

RFC 3165, *Definitions of Managed Objects for the Delegation of Management Scripts*

RFC 3231, *Definitions of Managed Objects for Scheduling Management Operations*

RFC 3273, *Remote Network Monitoring Management Information Base for High Capacity Networks*

RFC 3416, *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*

RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4)*

RFC 3419, *Textual Conventions for Transport Addresses*

RFC 3498, *Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures*

RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework*

RFC 3592, *Definitions of Managed Objects for the Synchronous Optical Network/ Synchronous Digital Hierarchy (SONET/SDH) Interface Type*

RFC 3593, *Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals*

RFC 3635, *Definitions of Managed Objects for the Ethernet-like Interface Types*

RFC 3637, *Definitions of Managed Objects for the Ethernet WAN Interface Sublayer*

RFC 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*

RFC 3877, *Alarm Management Information Base (MIB)*

RFC 3895, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*

RFC 3896, *Definitions of Managed Objects for the DS3/E3 Interface Type*

RFC 4001, *Textual Conventions for Internet Network Addresses*

RFC 4022, *Management Information Base for the Transmission Control Protocol (TCP)*

RFC 4113, *Management Information Base for the User Datagram Protocol (UDP)*

RFC 4220, *Traffic Engineering Link Management Information Base*
RFC 4273, *Definitions of Managed Objects for BGP-4*
RFC 4292, *IP Forwarding Table MIB*
RFC 4293, *Management Information Base for the Internet Protocol (IP)*
RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*
RFC 4511, *Lightweight Directory Access Protocol (LDAP): The Protocol*
RFC 4513, *Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms (TLS)*
RFC 4631, *Link Management Protocol (LMP) Management Information Base (MIB)*
RFC 4878, *Definitions and Managed Objects for Operations, Administration, and Maintenance (OAM) Functions on Ethernet-Like Interfaces*
RFC 5101, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information*
RFC 5102, *Information Model for IP Flow Information Export*
RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2 (TLS client, RSA public key)*
RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*
RFC 6425, *Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping*
RFC 6991, *Common YANG Data Types*
RFC 7420, *Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module*
RFC 7950, *The YANG 1.1 Data Modeling Language*
SFLOW-MIB, *sFlow MIB Version 1.3 (Draft 5)*

Multiprotocol Label Switching — Transport Profile (MPLS-TP)

RFC 5586, *MPLS Generic Associated Channel*
RFC 5921, *A Framework for MPLS in Transport Networks*
RFC 5960, *MPLS Transport Profile Data Plane Architecture*
RFC 6370, *MPLS Transport Profile (MPLS-TP) Identifiers*
RFC 6378, *MPLS Transport Profile (MPLS-TP) Linear Protection*
RFC 6426, *MPLS On-Demand Connectivity and Route Tracing*
RFC 6427, *MPLS Fault Management Operations, Administration, and Maintenance (OAM)*
RFC 6428, *Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile*
RFC 6478, *Pseudowire Status for Static Pseudowires*

RFC 7213, *MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing*

Multiprotocol Label Switching (MPLS)

draft-ietf-teas-sr-rsvp-coexistence-rec-02, *Recommendations for RSVP-TE and Segment Routing LSP co-existence*

RFC 3031, *Multiprotocol Label Switching Architecture*

RFC 3032, *MPLS Label Stack Encoding*

RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*

RFC 4023, *Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)*

RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*

RFC 5332, *MPLS Multicast Encapsulations*

RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*

RFC 6790, *The Use of Entropy Labels in MPLS Forwarding*

RFC 7510, *Encapsulating MPLS in UDP*

Network Address Translation (NAT)

draft-ietf-behave-address-format-10, *IPv6 Addressing of IPv4/IPv6 Translators*

draft-ietf-behave-v6v4-xlate-23, *IP/ICMP Translation Algorithm*

draft-miles-behave-l2nat-00, *Layer2-Aware NAT*

draft-nishitani-cgn-02, *Common Functions of Large Scale NAT (LSN)*

RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*

RFC 5382, *NAT Behavioral Requirements for TCP*

RFC 5508, *NAT Behavioral Requirements for ICMP*

RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*

RFC 6333, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*

RFC 6334, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite*

RFC 6887, *Port Control Protocol (PCP)*

RFC 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*

RFC 7915, *IP/ICMP Translation Algorithm*

Network Configuration Protocol (NETCONF)

RFC 5277, *NETCONF Event Notifications*

RFC 6022, *YANG Module for NETCONF Monitoring*
RFC 6241, *Network Configuration Protocol (NETCONF)*
RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*
RFC 6243, *With-defaults Capability for NETCONF*
RFC 7895, *YANG Module Library*

Open Shortest Path First (OSPF)

draft-ietf-ospf-ospfv3-lsa-extend-13, *OSPFv3 LSA Extendibility*
RFC 1586, *Guidelines for Running OSPF Over Frame Relay Networks*
RFC 1765, *OSPF Database Overflow*
RFC 2328, *OSPF Version 2*
RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*
RFC 3509, *Alternative Implementations of OSPF Area Border Routers*
RFC 3623, *Graceful OSPF Restart Graceful OSPF Restart (helper mode)*
RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*
RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*
RFC 4222, *Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance*
RFC 4552, *Authentication/Confidentiality for OSPFv3*
RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*
RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*
RFC 5185, *OSPF Multi-Area Adjacency*
RFC 5187, *OSPFv3 Graceful Restart (helper mode)*
RFC 5243, *OSPF Database Exchange Summary List Optimization*
RFC 5250, *The OSPF Opaque LSA Option*
RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*
RFC 5340, *OSPF for IPv6*
RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*
RFC 5838, *Support of Address Families in OSPFv3*
RFC 6987, *OSPF Stub Router Advertisement*
RFC 7684, *OSPFv2 Prefix/Link Attribute Advertisement*
RFC 7770, *Extensions to OSPF for Advertising Optional Router Capabilities*

OpenFlow

TS-007, *OpenFlow Switch Specification Version 1.3.1* (OpenFlow-hybrid switches)

Path Computation Element Protocol (PCEP)

draft-alvarez-pce-path-profiles-04, *PCE Path Profiles*

draft-ietf-pce-segment-routing-08, *PCEP Extensions for Segment Routing*

draft-ietf-pce-stateful-pce-14, *PCEP Extensions for Stateful PCE*

RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*

Point-to-Point Protocol (PPP)

RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*

RFC 1377, *The PPP OSI Network Layer Control Protocol (OSINLCP)*

RFC 1661, *The Point-to-Point Protocol (PPP)*

RFC 1662, *PPP in HDLC-like Framing*

RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*

RFC 1989, *PPP Link Quality Monitoring*

RFC 1990, *The PPP Multilink Protocol (MP)*

RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*

RFC 2153, *PPP Vendor Extensions*

RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*

RFC 2615, *PPP over SONET/SDH*

RFC 2686, *The Multi-Class Extension to Multi-Link PPP*

RFC 2878, *PPP Bridging Control Protocol (BCP)*

RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*

RFC 5072, *IP Version 6 over PPP*

Policy Management and Credit Control

3GPP TS 29.212 Release 11, *Policy and Charging Control (PCC); Reference points (Gx support as it applies to wireline environment (BNG))*

RFC 3588, *Diameter Base Protocol*

RFC 4006, *Diameter Credit-Control Application*

Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS Interworking*
MFA Forum 12.0.0, *Multiservice Interworking - Ethernet over MPLS*
MFA Forum 13.0.0, *Fault Management for Multiservice Interworking v1.0*
MFA Forum 16.0.0, *Multiservice Interworking - IP over MPLS*
MFA Forum 9.0.0, *The Use of Virtual trunks for ATM/MPLS Control Plane Interworking*
RFC 3916, *Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)*
RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3)*
RFC 4385, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*
RFC 4446, *IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)*
RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*
RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*
RFC 4619, *Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks*
RFC 4717, *Encapsulation Methods for Transport Asynchronous Transfer Mode (ATM) over MPLS Networks*
RFC 4816, *Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service*
RFC 5085, *Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires*
RFC 5659, *An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge*
RFC 5885, *Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)*
RFC 6073, *Segmented Pseudowire*
RFC 6310, *Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping*
RFC 6391, *Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network*
RFC 6575, *Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs*
RFC 6718, *Pseudowire Redundancy*
RFC 6829, *Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6*
RFC 6870, *Pseudowire Preferential Forwarding Status bit*
RFC 7023, *MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking*

RFC 7267, *Dynamic Placement of Multi-Segment Pseudowires*

Quality of Service (QoS)

RFC 2430, *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*

RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*

RFC 2598, *An Expedited Forwarding PHB*

RFC 3140, *Per Hop Behavior Identification Codes*

RFC 3260, *New Terminology and Clarifications for Diffserv*

Remote Authentication Dial In User Service (RADIUS)

RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*

RFC 2866, *RADIUS Accounting*

RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*

RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*

RFC 2869, *RADIUS Extensions*

RFC 3162, *RADIUS and IPv6*

RFC 4818, *RADIUS Delegated-IPv6-Prefix Attribute*

RFC 5176, *Dynamic Authorization Extensions to RADIUS*

RFC 6911, *RADIUS attributes for IPv6 Access Networks*

RFC 6929, *Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions*

Resource Reservation Protocol — Traffic Engineering (RSVP-TE)

draft-newton-mpls-te-dynamic-overbooking-00, *A Diffserv-TE Implementation Model to dynamically change booking factors during failure events*

RFC 2702, *Requirements for Traffic Engineering over MPLS*

RFC 2747, *RSVP Cryptographic Authentication*

RFC 2961, *RSVP Refresh Overhead Reduction Extensions*

RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value*

RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*

RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions (IF_ID RSVP_HOP object with unnumbered interfaces and RSVP-TE graceful restart helper procedures)*

RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*

RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*

RFC 3906, *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels*

RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*

RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*

RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*

RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*

RFC 5151, *Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*

RFC 5712, *MPLS Traffic Engineering Soft Preemption*

RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*

Routing Information Protocol (RIP)

RFC 1058, *Routing Information Protocol*

RFC 2080, *RIPng for IPv6*

RFC 2082, *RIP-2 MD5 Authentication*

RFC 2453, *RIP Version 2*

Segment Routing (SR)

draft-filsfils-spring-segment-routing-policy-05, *Segment Routing Policy for Traffic Engineering*

draft-francois-rtgwg-segment-routing-ti-lfa-04, *Topology Independent Fast Reroute using Segment Routing*

draft-gredler-idr-bgp-ls-segment-routing-ext-03, *BGP Link-State extensions for Segment Routing*

draft-ietf-idr-segment-routing-te-policy-02, *Advertising Segment Routing Policies in BGP*

draft-ietf-isis-segment-routing-extensions-04, *IS-IS Extensions for Segment Routing*

draft-ietf-mpls-spring-lsp-ping-02, Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing

draft-ietf-spring-conflict-resolution-05, Segment Routing MPLS Conflict Resolution

draft-ietf-spring-segment-routing-ldp-interop-09, Segment Routing interworking with LDP

Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH)

ANSI T1.105.03, Jitter Network Interfaces

ANSI T1.105.06, Physical Layer Specifications

ANSI T1.105.09, Network Timing and Synchronization

ITU-T G.703, Physical/electrical characteristics of hierarchical digital interfaces

ITU-T G.707, Network node interface for the synchronous digital hierarchy (SDH)

ITU-T G.813, Timing characteristics of SDH equipment slave clocks (SEC)

ITU-T G.823, The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy

ITU-T G.824, The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy

ITU-T G.825, The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)

ITU-T G.841, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1, issued in July 2002

ITU-T G.957, Optical interfaces for equipments and systems relating to the synchronous digital hierarchy

Time Division Multiplexing (TDM)

ANSI T1.403, DS1 Metallic Interface Specification

ANSI T1.404, DS3 Metallic Interface Specification

Timing

GR-1244-CORE, Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005

GR-253-CORE, SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000

IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*

ITU-T G.781, *Synchronization layer functions*, issued 09/2008

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC)*, issued 03/2003

ITU-T G.8261, *Timing and synchronization aspects in packet networks*, issued 04/2008

ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC)*, issued 08/2007

ITU-T G.8264, *Distribution of timing information through packet networks*, issued 10/2008

ITU-T G.8265.1, *Precision time protocol telecom profile for frequency synchronization*, issued 10/2010

ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network*, issued 07/2014

RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*

Two-Way Active Measurement Protocol (TWAMP)

RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP) (server, unauthenticated mode)*

RFC 5938, *Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)*

RFC 6038, *Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features*

Virtual Private LAN Service (VPLS)

RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*

RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*

RFC 5501, *Requirements for Multicast Support in Virtual Private LAN Services*

RFC 6074, *Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)*

RFC 7041, *Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging*

RFC 7117, *Multicast in Virtual Private LAN Service (VPLS)*

Voice and Video

DVB BlueBook A86, *Transport of MPEG-2 TS Based DVB Services over IP Based Networks*

ETSI TS 101 329-5 Annex E, *QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring*

ITU-T G.1020 Appendix I, *Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks - Mean Absolute Packet Delay Variation & Markov Models*

ITU-T G.107, *The E Model - A computational model for use in planning*

ITU-T P.564, *Conformance testing for voice over IP transmission quality assessment models*

RFC 3550 Appendix A.8, *RTP: A Transport Protocol for Real-Time Applications* (estimating the interarrival jitter)

RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*

RFC 4588, *RTP Retransmission Payload Format*

Wireless Local Area Network (WLAN) Gateway

3GPP TS 23.402, *Architecture enhancements for non-3GPP accesses* (S2a roaming based on GPRS)

Customer Document and Product Support



Customer Documentation

[Customer Documentation Welcome Page](#)



Technical Support

[Product Support Portal](#)



Documentation Feedback

[Customer Documentation Feedback](#)

