



Alcatel-Lucent 7950

EXTENSIBLE ROUTING SYSTEM | RELEASE 13.0.R4

LAYER 3 SERVICES GUIDE: INTERNET ENHANCED SERVICES AND VIRTUAL PRIVATE ROUTED NETWORK SERVICES

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

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

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

Copyright 2015 Alcatel-Lucent.

All rights reserved.

Disclaimers

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

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

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

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

Table of Contents

Preface	11
About This Guide	11
Audience	11
List of Technical Publications	12
Searching for Information	13
To search for specific information in this guide	13
To search for specific information in multiple documents	13
Technical Support	14
Internet Enhanced Service	
In This Chapter	15
IES Service Overview	16
IES Features	17
IP Interfaces	17
QoS Policy Propagation Using BGP (QPPB)	18
QPPB	21
QPPB and GRT Lookup	26
Object Grouping and State Monitoring	29
SAPs	31
Encapsulations	31
Pseudowire SAPs	32
Encapsulation	34
Pseudowire SAP Configuration	35
QoS for Pseudowire Ports and Pseudowire SAPs	36
Shaping and Bandwidth Control	37
Lag Considerations	38
Last Mile Packet Size Adjustment	38
Redundancy with Pseudowire SAPs	39
Operational Group Support for PW Ports	41
Routing Protocols	43
CPE Connectivity Check	43
QoS Policies	44
Filter Policies	44
Configuring an IES Service with CLI	45
Basic Configuration	46
Common Configuration Tasks	47
Configuring IES Components	48
Configuring an IES Service	48
Configuring IES Interface Parameters	49
Configuring SAP Parameters	50
IGMP Host Tracking	51
Service Management Tasks	52
Modifying IES Service Parameters	52
Deleting an IES Service	53
Disabling an IES Service	54

Table of Contents

Re-Enabling an IES Service	54
IES Services Command Reference	55
Command Hierarchies	55
IES Service Configuration Commands	71
IES Show Commands	173

Virtual Private Routed Network Service

In This Chapter	223
VPRN Service Overview	225
Routing Prerequisites	227
Core MP-BGP Support	228
Route Distinguishers	229
eiBGP Load Balancing	229
Route Reflector	232
CE to PE Route Exchange	233
Constrained Route Distribution (RT Constraint)	235
Constrained VPN Route Distribution Based on Route Targets	235
Configuring the Route Target Address Family	235
Originating RT Constraint Routes	235
Receiving and Re-Advertising RT Constraint Routes	236
Using RT Constraint Routes	237
BGP Fast Reroute in a VPRN	239
BGP Fast Reroute in a VPRN Configuration	240
VPRN Features	241
IP Interfaces	241
QoS Policy Propagation Using BGP (QPPB)	242
Object Grouping and State Monitoring	243
VPRN IP Interface Applicability	243
SAPs	245
Encapsulations	245
QoS Policies	246
Filter Policies	246
DSCP Marking	247
Default DSCP Mapping Table	249
Configuration of TTL Propagation for VPRN Routes	249
CE to PE Routing Protocols	251
PE to PE Tunneling Mechanisms	251
Per VRF Route Limiting	251
Spoke SDPs	252
T-LDP Status Signaling for Spoke-SDPs Terminating on IES/VPRN	253
IP-VPNs	254
Using OSPF in IP-VPNs	254
IPCP Subnet Negotiation	255
Cflowd for IP-VPNs	256
Carrier Supporting Carrier (CsC)	257
Terminology	257
CSC Connectivity Models	259
Traffic Leaking to GRT	260
Traffic Leaking from VPRN to GRT for IPv6	261

RIP Metric Propagation in VPRNs	262
NTP Within a VPRN Service	263
PTP Within a VPRN Service	263
VPN Route Label Allocation	264
QoS on Ingress Bindings	266
Multicast in IP-VPN Applications	268
Use of Data MDTs	270
Multicast Protocols Supported in the Provider Network	271
MVPN Membership Auto-discovery using BGP	271
MVPN (Rosen) Membership Auto-Discovery using BGP MDT-SAFI	274
PE-PE Transmission of C-Multicast Routing using BGP	274
VRF Route Import Extended Community	274
Provider Tunnel Support	276
Point-to-Multipoint Inclusive (I-PMSI) and Selective (S-PMSI) Provider Multicast Service Interface	276
Wildcard (C-*, C-*) P2MP LSP S-PMSI	276
P2MP LSP S-PMSI	279
Dynamic Multicast Signaling over P2MP LDP in VRF	279
MVPN Sender-only/Receiver-only	282
S-PMSI Trigger Thresholds	283
Migration from Existing Rosen Implementation	284
MVPN (NG-MVPN) Upstream Multicast Hop Fast Failover	285
Multicast VPN Extranet	286
Multicast Extranet for Rosen MVPN for PIM SSM	286
Multicast Extranet for NG-MVPN for PIM SSM	287
Multicast Extranet with Per-group Mapping for PIM SSM	288
Multicast GRT-source/VRF-receiver Extranet with Per Group Mapping for PIM SSM	290
Multicast Extranet with Per-group Mapping for PIM ASM	291
IPv6 MVPN Support	293
Multicast Core Diversity for Rosen MDT_SAFI MVPNs	296
NG-MVPN Multicast Source Geo-Redundancy	298
Multicast Core Diversity for Rosen MDT SAFI MVPNs	301
Inter-AS MVPN	303
BGP Connector Attribute	303
PIM RPF Vector	303
Inter-AS MVPN Option B	304
Inter-AS MVPN Option C	305
FIB Prioritization	307
Configuring a VPRN Service with CLI	309
Basic Configuration	310
Common Configuration Tasks	312
Configuring VPRN Components	313
Creating a VPRN Service	313
Configuring Global VPRN Parameters	314
Configuring VPRN Log Parameters	315
Configuring IPSec Parameters	325
Service Management Tasks	326
Modifying VPRN Service Parameters	326
Deleting a VPRN Service	327

Table of Contents

Disabling a VPRN Service	328
Re-enabling a VPRN Service	330
VPRN Services Command Reference	331
Command Hierarchies	331
VPRN Service Configuration Commands	377
VPRN Show Commands	699
Common CLI Command Descriptions	
In This Chapter	849
Common Service Commands	850
Standards and Protocol Support	853

List of Tables

Table 1: List of Technical Publications.	12
--	----

Internet Enhanced Service

Table 2: QPPB Interactions with SAP Ingress QoS.	27
Table 3: Packet Sizes Used for Pseudowire SAPs	39
Table 4: Default QinQ and TopQ SAP Dot1P Evaluation	146
Table 5: Top Position QinQ and TopQ SAP Dot1P Evaluation	147
Table 6: Bottom Position QinQ and TopQ SAP Dot1P Evaluation	147
Table 7: Default Dot1P Explicit Marking Actions	148
Table 8: QinQ Mark Top Only Explicit Marking Actions.	148

Virtual Private Routed Network Service

Table 9: BGP Fast Reroute Scenarios (VPRN Context)	239
Table 10: DSCP/FC Marking	247
Table 11: Label allocation logic	265
Table 12: Supported Configuration Combinations.	272
Table 13: Supported Configuration Combinations.	272
Table 14: Bottom Position QinQ and TopQ SAP Dot1P Evaluation	533
Table 15: Potential Adjacency Capabilities	651
Table 16: SNMP Trap Group Output Fields.	787
Table 17: Show Log Syslog Output Fields.	788

List of Tables

List of Figures

Internet Enhanced Service

Figure 1: Internet Enhanced Service	16
Figure 2: Use of QPPB to Differentiate Traffic in an ISP Network	20
Figure 3: Network Architecture using Pseudowire SAPs	33
Figure 4: PW SAP Egress Scheduling Hierarchy Options	38
Figure 5: Dual Homing into Multiple Layer 3 PEs	40
Figure 6: Master-Slave PW Redundancy	41

Virtual Private Routed Network Service

Figure 7: Virtual Private Routed Network	226
Figure 8: Route Distinguisher	229
Figure 9: Basic eiBGP Topology	230
Figure 10: Extranet Load Balancing	231
Figure 11: SDP-ID and VC Label Service Identifiers	252
Figure 12: Active/Standby VRF Using Resilient Layer 2 Circuits	253
Figure 13: CPEs Network Up-link Mode	255
Figure 14: Carrier Supporting Carrier Reference Diagram	257
Figure 15: RIP Metric Propagation in VPRNs	262
Figure 16: Ingress QoS Control on VPRN Bindings	266
Figure 17: Multicast in IP-VPN Applications	268
Figure 18: Dynamic mLDP Signaling for IP Multicast in VPRN	280
Figure 19: MVPN Sender-only/Receiver-only Example	283
Figure 20: Multicast VPN Traffic Flow	287
Figure 21: Source PE Transit Replication and Receiver PE Per-group SSM Extranet Mapping	289
Figure 22: GRT/VRF Extranet	291
Figure 23: Multicast Extranet with per group PIM ASM mapping	292
Figure 24: IPv6 MVPN Example	294
Figure 25: Multicast Core Diversity	296
Figure 26: Preferred Source Selection for Multicast Source Geo-Redundancy	298
Figure 27: Multicast Core Diversity	301
Figure 28: Inter-AS Option B Default MDT Setup	304
Figure 29: Inter-AS Option C Default MDT Setup	305
Figure 30: OSPF Areas	684

List of Figures

About This Guide

This guide describes Layer 3 service functionality provided by Alcatel-Lucent'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.

Audience

This guide is intended for network administrators who are responsible for configuring the 7950 XRS routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Concepts described in this guide include the following:

- Internet Enhanced Services (IES)
- Virtual Private Routed Network (VPRN) services

List of Technical Publications

The 7950 XRS documentation set is composed of the following guides:

Table 1: List of Technical Publications

Guide	Description
7950 XRS Basic System Configuration Guide	This guide describes basic system configurations and operations.
7950 XRS System Management Guide	This guide describes system security and access configurations as well as event logging and accounting logs.
7950 XRS Interface Configuration Guide	This guide describes XMA Control Module (XCM), XRS Media Adaptor (XMA), port and Link Aggregation Group (LAG) provisioning.
7950 XRS Router Configuration Guide	This guide describes logical IP routing interfaces and associated attributes such as an IP address, as well as IP and MAC-based filtering, and VRRP and Cflowd.
7950 XRS Routing Protocols Guide	This guide provides an overview of routing concepts and provides configuration examples for RIP, OSPF, IS-IS, BGP, and route policies.
7950 XRS MPLS Guide	This guide describes how to configure Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP).
7950 XRS Services Guide	This guide describes how to configure service parameters such as service distribution points (SDPs), customer information, and user services.
7950 XRS Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN	This guide describes Virtual Leased Lines (VLL), Virtual Private LAN Service (VPLS), Provider Backbone Bridging (PBB), and Ethernet VPN (EVPN).
7950 XRS Layer 3 Services Guide: Internet Enhanced Services and Virtual Private Routed Network Services	This guide describes Internet Enhanced Services (IES) and Virtual Private Routed Network (VPRN) services.

Table 1: List of Technical Publications

Guide	Description
7950 XRS OAM and Diagnostics Guide	This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.
7950 XRS Quality of Service Guide	This guide describes how to configure Quality of Service (QoS) policy management.

Searching for Information

You can use Adobe Reader, Release 6.0 or later, to search one or more PDF files for a term.

To search for specific information in this guide

1. From the Adobe Reader main menu, choose Edit > Search or Advanced Search. The Search panel opens.
2. Click on the In the current document radio button.
3. Enter the term to search for.
4. Select the following search criteria, if required:
 - Whole words only
 - Case-Sensitive
 - Include Bookmarks
 - Include Comments
5. Click on the Search button. Adobe Reader displays the search results.

You can expand the entries by clicking on the + symbol.

To search for specific information in multiple documents

Note: The PDF files that you search must be in the same folder.

1. From the Adobe Reader main menu, choose Edit > Search or Advanced Search. The Search panel opens.

2. Click on the All PDF Documents in radio button.
3. Choose the folder in which to search using the drop-down menu.
4. Enter the term to search for.
5. Select the following search criteria, if required:
 - Whole words only
 - Case-Sensitive
 - Include Bookmarks
 - Include Comments
6. Click on the Search button. Adobe Reader displays the search results.

You can expand the entries for each file by clicking on the + symbol.

Technical Support

If you purchased a service agreement for your 7950 SR-Series router and related products from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance. If you purchased an Alcatel-Lucent service agreement, follow this link to contact an Alcatel-Lucent support representative and to access product manuals and documentation updates:

<https://support2.alcatel-lucent.com/portal/olcsHome.do>

Internet Enhanced Service

In This Chapter

This chapter provides information about Internet Enhanced Service (IES), process overview, and implementation notes.

Topics in this chapter include:


- [IES Service Overview on page 16](#)
- [IES Features on page 17](#)
 - [IP Interfaces on page 17](#)
 - [QoS Policy Propagation Using BGP \(QPPB\) on page 18](#)
 - [Object Grouping and State Monitoring on page 29](#)
 - [Filter Policies on page 44](#)
- [Configuring an IES Service with CLI on page 45](#)
- [Basic Configuration on page 46](#)
- [Common Configuration Tasks on page 47](#)
- [Service Management Tasks on page 52](#)

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.



OSSG023

Figure 1: Internet Enhanced Service

IES Features

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

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 \(QPPB\) on page 18](#)
- ICMP Options

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

- MPLS forwarding
- NTP broadcast receipt

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 on page 15 and the IP Router Configuration section in the 7x50 OS 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 routers, 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, 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 ?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. 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.

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.
-

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.

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. Note however, that the DSCP or other COS markings could be left unchanged in the ISP's network and QPPB used on every node.



OSSG639

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

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.
- 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.

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 h1 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>vpn>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)

Note that 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 using the following modified versions of the static-route commands:

- `static-route {ip-prefix/prefix-length | ip-prefix netmask} [fc fc-name [priority {low | high}]] next-hop ip-int-name|ip-address`
- `static-route {ip-prefix/prefix-length | ip-prefix netmask} [fc fc-name [priority {low | high}]] indirect ip-address`

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.

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                               0
  h1, high
-----
No. of Routes: 1
=====
A:Dut-A#
```

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. Note however, 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)
- IES SAP and spoke SDP interfaces (config>service>ies>interface)

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 (see section 5.7 for further details). 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>vprn>nw-if).

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

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. In release 9.0 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.

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 on page 21](#) 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 on page 21](#) 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.

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

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.

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 subclass (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 2](#) summarizes these interactions.

Table 2: 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
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

Table 2: 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	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
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

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.

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 is example uses an R-VPLS context. The VPLS instance includes the **allow-ip-int-binding** and the **service-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 customer 1 create
  allow-ip-int-binding
  stp
    shutdown
  exit
  service-name "v1"
  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 21.1.1.1/24
    monitor-oper-group "g1"
    vpls "v1"
  exit
no shutdown
exit
```

SAPs

Encapsulations

The following SAP encapsulations are supported on the 7950 IES services:

- Ethernet null
- Ethernet dot1q

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 SR OS Triple Play Service Delivery Architecture user 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

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, 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 note that 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.

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:

1. Define a pseudowire port

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

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
```

3. Perform one of the following steps:

- a. 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 30.1.1.1/24
mac 00:00:00:00:00:ff
static-arp 30.1.1.2 00:00:00:00:00:aa
sap pw-1:1 create
exit
exit
no shutdown
exit

```

- b. 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

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 IOM3 with Ethernet MDA or HSMDAv2, and policy based redirection for the IOM3 with Ethernet MDA, as applicable.
- Policing and H-POL

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 will inherit the vport used by the pw-port egress shaping.

Note that 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(s) 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 (**configure>port>ethernet>access>egress>vport vport-name create**).
2. Explicitly assign (via 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 'hsmda-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 SR OS Quality of Service guide for more details.




Figure 4: PW SAP Egress Scheduling Hierarchy Options

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.

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 encaps 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.

Note that 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 3](#) summarizes the default packet sizes used at each of the schedulers on the IOM/Ethernet MDA and HSMDAv2, assuming a 1000byte customer packet.

Table 3: Packet Sizes Used for Pseudowire SAPs

Type	Size
exp-secondary-shaper	20B preamble + 26 MPLS + 1000B pkt
queue/policer rate on hsmdav2	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
vport scheduler rate	1000B pkt
vport scheduler to port-scheduler rates	20B preamble + 1000B pkt

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/VPRN.

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 e.g. 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.

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). Note that, 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 i.e. 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"

```

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 that 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.

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 ICPM 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.

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. Note that both Layer 2 or Layer 3 criteria can be used in the QoS policies for traffic classification in an IES.

Filter Policies

Only IP filter policies can be applied to IES services.

Configuring an IES Service with CLI

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

Topics in this section include:

- [Basic Configuration on page 46](#)
- [Common Configuration Tasks on page 47](#)
 - [Configuring IES Components on page 48](#)
 - [Configuring an IES Service on page 48](#)
 - [Configuring IES Interface Parameters on page 49](#)
 - [Configuring SAP Parameters on page 50](#)
- [Service Management Tasks on page 52](#)
 - [Modifying IES Service Parameters on page 52](#)
 - [Deleting an IES Service on page 53](#)
 - [Disabling an IES Service on page 54](#)
 - [Re-Enabling an IES Service on page 54](#)

Basic Configuration

The most basic IES service configuration has the following entities:

- Customer ID (refer to the *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 example displays 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#
```

Common Configuration Tasks

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

1. Associate an IES service with a customer ID.
2. Associate customer ID with the service.
3. Assign an IP address.
4. Create an interface.
5. Define SAP parameters on the interface
 - Select node(s) and port(s).
 - Optional — select filter policies (configured in the **config>filter** context).
6. Enable service.

Configuring IES Components

Use the CLI syntax to configure the following entities:

- [Configuring an IES Service on page 48](#)
 - [Configuring IES Interface Parameters on page 49](#)
 - [Configuring SAP Parameters on page 50](#)
-

Configuring an IES Service

Use the following CLI syntax to create an IES service:

The following example displays a basic IES service configuration.

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

Configuring IES Interface Parameters

The following example displays an IES configuration with interface parameters:

```
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#
```

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.

Filter policies are configured in the **config>filter** context and must be explicitly applied to a SAP. There are no default filter policies.

This example displays an IES SAP configuration.

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

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
            delayed-enable 1200
            address 64.64.64.64/24
            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 2.3.4.5 255.255.255.0
            secondary 2.3.4.5/24
            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#
```

Service Management Tasks

This section discusses the following service management tasks:

- [Modifying IES Service Parameters on page 52](#)
 - [Deleting an IES Service on page 53](#)
-

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 displays the modified service:

```
*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#
```

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
```

Disabling an IES Service

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

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

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`

IES Services Command Reference

Command Hierarchies

- [IES Service Configuration Commands on page 55](#)
- [Global Commands on page 55](#)
- [Interface Commands on page 56](#)
- [Interface SAP Commands on page 59](#)
- [VRRP Commands on page 63](#)
- [Spoke SDP Commands on page 65](#)
- [Show Commands on page 67](#)

IES Service Configuration Commands

Global Commands

```

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

```

Interface Commands

```

config
  — service
    — ies service-id [customer customer-id] [vpn vpn-id]
      — 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
        — arp-limit limit [log-only] [threshold percent]
        — no arp-limit
        — arp-retry-timer timer-multiple
        — no arp-retry-timer
        — [no] arp-timeout
        — 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
        — no cpu-protection
        — delayed-enable seconds
        — no delayed-enable
        — description description-string
        — no description
          — policy-name
          — no
        — dynamic-tunnel-redundant-next-hop ip-address
        — no dynamic-tunnel-redundant-next-hop
        — [no] enable-ingress-stats
        — [no] enable-mac-accounting
        — [no] flowspec
        — host-connectivity-verify [source {vrrp | interface}] [interval interval]
          [action {remove | alarm}]
        — icmp
          — [no] mask-reply
          — redirects [number seconds]
          — no redirects
          — ttl-expired [number seconds]
          — no ttl-expired
          — unreachables [number seconds]
          — no unreachables
        — if-attribute
          — [no] admin-group group-name [group-name...(up to 5 max)]
          — no admin-group
          — [no] srlg-group group-name [group-name...(up to 5 max)]
          — no srlg-group
        — ip-load-balancing {source | destination}
        — no ip-load-balancing
        — 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**]
 - **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**
 - **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* [**preferred**]
 - **no link-local-address**
 - **[no] local-proxy-nd**
 - **neighbor** *ipv6-address mac-address*
 - **no neighbor** *ipv6-address*
 - **neighbor-limit**
 - **no neighbor-limit** *limit* [**log-only**] [**threshold percent**]
 - **proxy-nd-policy** *policy-name* [*policy-name...*(up to 5 max)]
 - **no proxy-nd-policy**
 - **[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**}
 - **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**
- **[no] local-proxy-arp**
- **[no] loopback**
- **[no] mac** *ieee-address*
- **monitor-oper-group** *name*
- **no monitor-oper-group**
- **multicast-network-domain** *multicast-network-domain*
- **no multicast-network-domain**
- **[no] proxy-arp-policy** *policy-name* [*policy-name...*(up to 5 max)]
- **[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*
- **[no] shutdown**
- **static-arp** *ieee-mac-addr* *unnumbered*
- **no static-arp** *unnumbered*
- **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**

Interface SAP Commands

```

config
  — service
    — ies service-id [customer customer-id] [vpn vpn-id]
      — [no] interface ip-int-name
        — [no] sap sap-id
          — arp aarpId type type
          — no arp
          — accounting-policy acct-policy-id
          — no accounting-policy [acct-policy-id]
          — anti-spoof {ip | ip-mac}
          — no anti-spoof
          — app-profile app-profile-name
          — no app-profile
          — calling-station-id calling-station-id
          — no calling-station-id
          — description description-string
          — no description
          — egress
            — agg-rate-limit agg-rate [queue-frame-based-accounting]
            — no agg-rate-limit
            — [no] agg-rate
              — rate {max | rate}
              — no rate
              — [no] limit-unused-bandwidth
              — [no] queue-frame-based-accounting
            — [no] qinq-mark-top-only
            — filter [ip ip-filter-id]
            — filter [ipv6 ipv6-filter-id]
            — no filter [ip ip-filter-id] [ipv6 ipv6-filter-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-id]
            — no qos
            — [no] queue-override
              — [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
                — high-prio-only percent
                — no high-prio-only
                — mbs size-in-kbytes
                — 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**
- **eth-cfm**
 - **[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**
 - **[no] ccm-padding-size** *ccm-padding*
 - **[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*
 - **[no] shutdown**
 - **[no] squelch-ingress-levels** [*md-level* [*md-level...*]]
 - **tunnel-fault** [**accept** | **ignore**]
- **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*]
 - **[no] flowspec**
 - **flowspec-ipv6**
 - **[no] flowspec-ipv6**
 - **match-qinq-dot1p** {**top** | **bottom**}
 - **no match-qinq-dot1p**
 - **qos** *policy-id* [**shared-queuing** | **multipoint-shared**][**fp-redirect-group** *queue-group-name* *instance-id*]
 - **no qos** *policy-id*
 - **[no] queue-override**
 - **[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**
- **high-prio-only** *percent*
- **no high-prio-only**
- **mbs** *size-in-kbytes*
- **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**
- **lag-link-map-profile** *lag-link-map-profile-id*
- **no lag-link-map-profile**
- **lag-per-link-hash** class {1 | 2 | 3} weight [1..1024]
- **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**
 - **[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**
- **tod-suite** *tod-suite-name*
- **no tod-suite**
- **transit-policy** *ip-aasub-policy-id*
- **no transit-policy**

— [no] **shutdown**

VRRP Commands

```

config
— service
    — ies service-id [customer customer-id] [vpn vpn-id]
        — [no] interface ip-int-name
            — [no] ipv6
                — vrrp virtual-router-id [owner]
                — no vrrp virtual-router-id
                    — [no] backup ip-address
                    — [no] bfd-enable service-id interface interface-name dst-ip ip-address
                    — [no] bfd-enable interface interface-name dst-ip ip-address
                    — init-delay seconds
                    — no init-delay
                    — mac mac-address
                    — no mac
                    — [no] master-int-inherit
                    — message-interval {[seconds] [milliseconds milliseconds]}
                    — no message-interval
                    — [no] ping-reply
                    — policy vrrp-policy-id
                    — no policy
                    — [no] preempt
                    — priority base-priority
                    — no priority
                    — [no] shutdown
                    — [no] standby-forwarding
                    — [no] telnet-reply
                    — [no] traceroute-reply
            — vrrp virtual-router-id [owner]
            — no vrrp virtual-router-id
                — authentication-key {authentication-key | hash-key} [hash | hash2]
                — no authentication-key
                — [no] backup ip-address
                — [no] bfd-enable [service-id] interface interface-name dst-ip ip-address
                — init-delay seconds
                — no init-delay
                — mac ieee-address
                — no mac
                — [no] master-int-inherit
                — message-interval {[seconds] [milliseconds milliseconds]}
                — no message-interval
                — [no] ping-reply
                — policy vrrp-policy-id
                — no policy
                — [no] preempt

```

- **priority** *priority*
- **no priority**
- **[no] shutdown**
- **[no] ssh-reply**
- **[no] standby-forwarding**
- **[no] telnet-reply**
- **[no] traceroute-reply**

Spoke SDP Commands

```

config
  — service
    — ies service-id [customer customer-id] [vpn vpn-id]
      — [no] interface ip-int-name
        — [no] spoke-sdp sdp-id:vc-id [vc-type {ether | ipipe}] [create]
          — arp arpId type type
          — no arp
          — 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
          — eth-cfm
            — [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
              — [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}
              — low-priority-defect seconds
              — [no] shutdown
            — [no] squelch-ingress-levels [md-level [md-level...]]
          — 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
            — vc-label egress-vc-label
            — no vc-label [egress-vc-label]
          — flowspec
          — [no] flowspec
          — flowspec-ipv6
          — [no] flowspec-ipv6

```

- **[no] hash-label**
- **ingress**
 - **filter** {**ip** *ip-filter-id*}
 - **filter** [**ipv6** *ipv6-filter-id*]
 - **no filter**
 - **[no] flowspec**
 - **qos** *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-aasub-policy-id*
- **no transit-policy**

Show Commands

```

show
  — service
    — customer [customer-id] [site customer-site-name]
    — sap-using [sap sap-id]
    — sap-using interface [ip-address | ip-int-name]
    — sap-using [ingress | egress] atm-td-profile id-profile-id
    — sap-using [ingress | egress] filter filter-id
    — sap-using [ingress | egress] qos-policy qos-policy-id
    — sap-using authentication-policy policy-name
    — service-using [ies] [customer customer-id]
    — sdp
      — 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
    — id service-id
      — all
      — arp [ip-address][mac ieee-address][sap sap-id][interface ip-int-name] [sdp sdp-id:vc-id]
      — base
      — gsm
        — neighbors group [name] [ip-address]
        — sessions [group name] neighbor ip-address [port port-number] [association] [statistics]
      — host
      — host-connectivity-verify statistics [sap sap-id]
      — interface [ip-address | ip-int-name] [interface-type] [detail] [family]
      — interface
      — retailers
      — sap sap-id [detail]
      — sap sap-id queue-depth [queue queue-id][ingress|egress]
      — sdp [{sdp-id | far-end ip-address}] [detail]
    — pw-port
      — pw-port [pw-port-id] [detail]
      — pw-port sdp [sdp-id]
      — pw-port sdp none
    — router
      — vrrp
        — instance
        — instance interface interface-name [vrid virtual-router-id]
        — instance interface interface-name vrid virtual-router-id ipv6
        — statistics

```

Clear Commands

```
clear
```

Command Hierarchies

- **router**
 - **interface** [*ip-int-name* | *ip-address*] [**icmp**]
 - **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**
 - **arp-host** { **mac** *ieee-address* | **sap** *sap-id* | **ip-address** *ip-address*[/*mask*] }
 - **arp-host** [**port** *port-id*] [**inter-dest-id** *intermediate-destination-id* | **no-inter-dest-id**]
 - **arp-host statistics** [**sap** *sap-id* | **interface** *interface-name*]
 - **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**]
 - **1 — 4294967295site** *service-id*
 - **spoke-sdp** *sdp-id*[:*vc-id*] *ingress-vc-label*
 - **stp**

Debug Commands

```

debug
  — service
    — id service-id
    — [no] arp-host
      — [no] host-connectivity-verify
        — [no] ip ip-address
        — [no] mac ieee-address
        — [no] sap sap-id
    — router
      — vrrp
        — [no] events
        — events interface-name [vrid virtual-router-id]
        — events interface-name vrid virtual-router-id ipv6
        — [no] packets
        — packets interface-name [vrid virtual-router-id]
        — packets interface-name vrid virtual-router-id ipv6

```

Monitor Commands

```

monitor
  — router
    — vrrp
      — instance interface interface-name vr-id virtual-router-id [ipv6] [interval seconds]
        [repeat repeat] [absolute| rate]

```

IES Service Configuration Commands

Generic Commands

shutdown

Syntax [no] shutdown

Context config>service>ies
 config>service>ies>igmp-snooping
 config>service>ies>if>sap>eth-cfm
 config>service>ies>if
 config>service>ies>if>vrrp

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.

Special Cases **IES** — 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.

For example if:

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

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

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

description

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

Context config>service>ies
 config>service>ies>if>dhcp

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

Generic Commands

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

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

Default No description associated with the configuration context.

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

IES Global Commands

ies

Syntax `ies service-id customer customer-id [vpn vpn-id] [create]`
`no ies service-id`

Context config>service

Description 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.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every SR OS router on which this service is defined.

Values	<i>service-id:</i>	1 — 2147483648
	<i>svc-name:</i>	64 characters maximum

customer *customer-id* — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 — 2147483647

vpn *vpn-id* — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.

Values 1 — 2147483647

Default null (0)

service-name

Syntax **service-name** *service-name*
no service-name

Context config>service>ies

Description This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a given service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the SR OS platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a given service once it is initially created.

Parameters *service-name* — Specifies a unique service name to identify the service. Service names may not begin with an integer (0-9).

igmp-host-tracking

Syntax **igmp-host-tracking**

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

Description This command enables the context to configure IGMP host tracking parameters.

expiry-time

Syntax **expiry-time** *expiry-time*
no expiry-time

Context config>service>ies>igmp-snooping
config>service>ies>sub-if>grp-if>sap>igmp-snooping

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 *expiry-time* — Specifies the time, in seconds, that this system continues to track an inactive host.

Values 1 — 65535

IES Interface Commands

interface

Syntax **interface** *ip-int-name* [**create**] [**tunnel**]
no interface *ip-int-name*

Context config>service>ies

Description This command creates a logical IP routing interface for an Internet Enhanced Service (IES). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.

The **interface** command, under the context of services, 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 IESservices, 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, etc.), 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.

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

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.

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.

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.

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. Note that 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. Note that 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.

delayed-enable

Syntax **delayed-enable** *seconds*
 no delayed-enable

Context f

Description This command will cause a delay in the activation of an IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface up.

IES Interface Commands

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 expires.

Parameters *seconds* — Specifies a delay, in seconds, to make the interface operational.

Values 1 — 1200

allow-directed-broadcasts

Syntax **[no] allow-directed-broadcasts**

Context config>service>ies>if

Description This command enables the forwarding of directed broadcasts out of the IP interface.

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

anti-spoof

Syntax **anti-spoof {ip | mac | ip-mac}**
no anti-spoof

Context config>service>ies>if>sap

Description This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP.

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**, **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 disables anti-spoof filtering on the SAP.

Default no anti-spoof

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 will fail.

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 will fail. 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 will fail. 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.

app-profile

Syntax **app-profile** *app-profile-name*
no app-profile

Context config>service>ies>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>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**.

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.

percent — The threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 — 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 ARP learning is disabled and no dynamic ARP entries are learned.

Values 0 — 524288

arp-retry-timer

Syntax	arp-retry-timer <i>timer-multiple</i> no arp-retry-timer
Context	config>service>ies>if
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 seconds.
Default	5 seconds
Parameters	<i>timer-multiple</i> — specifies the multiple of 100 ms that the ARP retry timer will be configured as. Values 1 — 300 (equally a timer range of 100 ms to 30,000 ms)

arp-timeout

Syntax	arp-timeout <i>seconds</i> no arp-timeout
Context	config>service>ies>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	14400 seconds
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 — 65535

bfd

Syntax	bfd <i>transmit-interval</i> [receive <i>receive-interval</i>] [multiplier <i>multiplier</i> [echo-receive <i>echo-interval</i>]] no bfd
Context	config>service>ies>if config>service>ies>if>ipv6
	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 interface.

Default no bfd

Parameters *transmit-interval* — Sets the transmit interval for the BFD session.

Values 100 — 100000

Default 100

receive receive-interval — Sets the receive interval for the BFD session.

Values 100 — 100000

Default 100

multiplier *multiplier* — Set the multiplier for the BFD session.

Values 3— 20

Default 3

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

Values 100 — 100000

Default 100

cflowd-parameters

Syntax **cflowd-parameters**
no cflowd-parameters

Context config>service>ies>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

IES Interface Commands

Syntax	sampling {unicast multicast} type {acl interface} [direction {ingress-only egress-only both}] no sampling {unicast multicast}
Context	config>service>ies>if>cflowd-parameters
Description	<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>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.</p> <p>The no form of the command disables the associated type of traffic sampling on the associated interface.</p>
Default	no sampling
Parameters	<p>unicast — Specifies that the sampling command will control the sampling of unicast traffic on the associated interface/SAP.</p> <p>multicast — Specifies that the sampling command will control the sampling of multicast traffic on the associated interface/SAP.</p> <p>type — Specifies the Cflowd sampling type on the given virtual router interfaces.</p> <p>Values</p> <ul style="list-style-type: none">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. <p>direction — Specifies the direction to collect traffic flow samples.</p> <p>Values</p> <ul style="list-style-type: none">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>ies>if
Description	<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 config>sys>security>cpu-protection>policy <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) cpu-protection 255 (for network interfaces)</p>

none (for video-interfaces, shown as no cpu-protection in CLI)

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 — 255

local-proxy-arp

Syntax [no] local-proxy-arp

Context config>service>ies>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

ip-mtu

Syntax ip-mtu *octets*
no ip-mtu

Context config>service>ies>if

Description This command configures the IP maximum transmit unit (packet) for this interface.

Note that because this connects a Layer 2 to a Layer 3 service, this parameter can be adjusted under the IES interface.

The MTU that is advertised from the IES size is:

$$\text{MINIMUM}((\text{SdpOperPathMtu} - \text{EtherHeaderSize}), (\text{Configured ip-mtu}))$$

By default (for ethernet network interface) if no ip-mtu is configured it is $(1568 - 14) = 1554$.

The **no** form of the command returns the default value.

Default no ip-mtu

ip-load-balancing

Syntax ip-load-balancing {source | destination}
no ip-load-balancing

Context config>service>ies>if

Description This command specifies whether to include source address or destination address or both in LAG/ECMP hash on IP interfaces. Additionally, when l4-load-balancing is enabled the command applies also to inclusion of source/destination port in the hash inputs.

IES Interface Commands

The **no** form of this command includes both source and destination parameters.

Default no ip-load-balancing

Parameters **source** — Specifies to use source address and (if l4-load balancing is enabled) source port in the hash, ignore destination address/port.
destination — Specifies to use destination address and (if l4-load balancing is enabled) destination port in the hash, ignore source address/port.

ipcp

Syntax ipcp

Context config>service>ies>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.

Default none

dns

Syntax dns ip-address [secondary ip-address]
dns secondary ip-address
no dns [ip-address] [secondary ip-address]

Context config>service>ies>if>ipcp

Description This command defines the dns address(es) to be assigned to the far-end of the associated PPP/MLPPP link through 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 *ip-address* — This parameter 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 — This parameter 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 <i>ip-address</i> no peer-ip-address
Context	config>service>ies>if>ipcp
Description	This command defines the remote IP address to be assigned to the far-end of the associated PPP/MLPPP link via IPCP extensions. This command is only applicable if the associated SAP/port is a PPP/MLPPP interface with an IPCP encapsulation. The no form of the command deletes the IPCP extension peer-ip-address configuration.
Default	no peer-ip-address (0.0.0.0)
Parameters	<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.

load-balancing

Syntax	load-balancing
Context	config>service>ies>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 { source destination inner-ip } no egr-ip-load-balancing
Context	config>service>ies>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 l4-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	source — Specifies using the source address and, if l4-load balancing is enabled, the source port in the hash, ignore destination address/port. destination — Specifies using the destination address and, if l4-load balancing is enabled, the destination port in the hash, ignore source address/port.

IES Interface Commands

inner-ip — 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

Syntax	[no] spi-load-balancing
Context	config>service>ies>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	disabled

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	disabled

loopback

Syntax	[no] loopback
Context	config>service>ies>if
Description	This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated IES interface cannot be bound to a SAP. Note that you can configure an IES interface as a loopback interface by issuing the loopback command instead of the sap command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.
Default	none

lag-per-link-hash

Syntax	lag-per-link-hash class {1 2 3} weight [1..1024] no 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)

mac

- Syntax** **mac** *ieee-address*
no mac
- Context** config>service>ies>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 <i>multicast-network-domain</i> 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

Syntax	secondary { <i>ip-address/mask</i> <i>ip-address netmask</i> } [broadcast all-ones host-ones] [igp-inhibit] no secondary <i>ip-address</i>
Context	config>service>ies>if
Description	This command assigns a secondary IP address/IP subnet/broadcast address format to the interface.
Default	none
Parameters	<p><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).</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. Note that 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>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.</p> <p>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 (all-ones) or the valid subnet broadcast address (host-ones) will be received by the IP interface. (<i>Default: host-ones</i>)</p>

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.

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.

static-arp

Syntax **static-arp** *ieee-mac-address unnumbered*
no static-arp *unnumbered*

Context config>service>ies>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 appears in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.

If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

The **no** form of the command removes a static ARP entry.

Default None

Parameters *ip-address* — Specifies the IP address for the static ARP in IP address dotted decimal notation.

ieee-mac-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

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** *ip-address*
no static-tunnel-redundant-next-hop

Context config>service>ies>if

Description This command specifies redundant next-hop address on public or private IPsec interface (with public or

IES Interface Commands

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.

Default none

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

tos-marking-state

Syntax **tos-marking-state** {**trusted** | **untrusted**}
no tos-marking-state

Context config>service>ies>if

Description 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.

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 **untrusted** for config>service>ies context

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.

unnumbered

Syntax	unnumbered [<i>ip-int-name</i> <i>ip-address</i>] no unnumbered
Context	config>service>ies>if
Description	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.
Parameters	<i>ip-int-name</i> — Specifies the name of an IP interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. <i>ip-address</i> — Specifies an IP address.

urpf-check

Syntax	[no] urpf-check
Context	config>service>ies>if config>service>ies>if>ipv6
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	disabled

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. The no form of the command removes 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 Value Added Service (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 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 Value Added Service (VAS) connectivity. For service chaining, traffic arriving from network interfaces (downstream) is redirected to a PBR target reachable over this interface for downstream VAS

IES Interface Commands

processing. Upstream traffic after VAS processing must arrive on this interface, so regular routing can be applied.

to-from-both — Used when a single interface is used for Value Added Service (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 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>ies>if>urfp-check
Description	This command specifies the mode of unicast RPF check. The no form of the command reverts to the default (strict) mode.
Default	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.

proxy-arp-policy

Syntax	[no] proxy-arp <i>policy-name</i> [<i>policy-name...</i> (up to 5 max)]
Context	config>service>ies>if
Description	This command configures a proxy ARP policy for the interface. The no form of this command disables the proxy ARP capability.
Default	no proxy-arp
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, etc.), the entire string must be enclosed within double quotes. The specified name(s) must already be defined.

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

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.

IES Interface Commands

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.

secure-nd

Syntax [no] **secure-nd**

Context config>service>ies>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>ies>if>ipv6>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** *modifier*
[no] **link-local-modifier**

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

Description 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** 512–1024

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** 0–1

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 - 65535

tcp-mss

Syntax	tcp-mss <i>mss-value</i> no tcp-mss
Context	config>service>ies>if config>service>ies>if>ipv6
Description	<p>This command statically sets the TCP maximum segment size (MSS) for TCP connections originated from the associated IP interface to the specified value.</p> <p>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 ($\text{tcp_mss} = \text{ip_mtu} - 40$).</p>
Default	no tcp-mss
Parameters	<p><i>mss-value</i> — The TCP MSS value that should be used in the TCP SYN packet during the three-way handshake negotiation of a TCP connection.</p> <p>Note: 9158 = max-IP_MTU (9198)-40</p> <p>Values 536 - 9158 (IPv4) 1220 - 9138 (IPv6)</p>

remote-proxy-arp

Context	config>service>ies>if
Description	<p>This command enables remote proxy ARP on the interface.</p> <p>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.</p>
Default	no remote-proxy-arp

arp-populate

Syntax	[no] arp-populate
Context	config>service>ies>if config>service>ies>sub-if>grp-if
Description	<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 lease-populate), and optionally with static entries entered with the host command.</p> <p>Enabling the arp-populate command will remove any dynamic ARP entries learned on this interface from the ARP cache.</p> <p>The arp-populate command will fail if an existing static ARP entry exists for this interface. The arp-</p>

populate command will fail 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 will fail.

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. The **arp-populate** command can only be enabled on IES and 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 for this interface will be removed from the system's ARP cache.

Default not enabled

host-lockout-policy

- Syntax** **host-lockout-policy** *policy-name*
no host-lockout-policy
- Context** config>service>ies>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>ies>if>sap
This command administratively enables host creation on this SAP.

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 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, etc.), the entire string must be enclosed within double quotes.

flowspec

- Syntax** [**no**] **flowspec**
- Context** config>service>ies>if>sap>ingress
- Description** This command enables flowspec filtering on an IP interface of the base router. Filtering is based on all of the flowspec routes that have been received and accepted by the base router. Ingress traffic on an IP interface can be filtered by both a user-defined ip filter and flowspec. In this case, the user-defined ip filter entries are evaluated before the flowspec routes and the default action of the user-defined ip filter applies as the very last rule.
The **no** form of the command removes flowspec filtering from an IP interface.

Default No interfaces have flowspec enabled.

IES Interface ICMP Commands

icmp

Syntax	icmp
Context	config>service>ies>if
Description	This command enables the context to configure Internet Control Message Protocol (ICMP) parameters on an IES service

mask-reply

Syntax	[no] mask-reply
Context	config>service>ies>if>icmp
Description	<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 mask-reply 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 no form of this command disables replies to ICMP mask requests on the router interface.</p>
Default	mask-reply — Reply to ICMP mask requests.

redirects

Syntax	redirects [<i>number seconds</i>] no redirects
Context	config>service>ies>if>icmp
Description	<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 redirects command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional <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>By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval. (<i>Default: redirects 100 10</i>)</p> <p>The no form of this command disables the generation of icmp redirects on the router interface.</p>

Default **redirects 100 10** — Maximum of 100 redirect messages in 10 seconds

Parameters *number* — The maximum number of ICMP redirect messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP redirect messages that can be issued.

Values 1 — 60

ttl-expired

Syntax **ttl-expired** *number seconds*
no ttl-expired

Context config>service>ies>if>icmp

Description This command configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default ttl-expired 100 10

Parameters *number* — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 — 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>ies>if>icmp

Description This command configures the rate for ICMP host and network destination unreachable messages issued on the router interface.

The **unreachables** command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional *number* and *time* parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval.

By default, generation of ICMP destination unreachable messages is enabled at a maximum rate of 10 per 60

IES Interface Commands

second time interval.

The **no** form of this command disables the generation of icmp destination unreachable messages on the router interface.

Default **unreachables 100 10**

Parameters *number* — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP unreachable messages that can be issued.

Values 1 — 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>interface>if-attribute
 config>service>ies>interface>if-attribute
 config>service>vprn>interface>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.

It should be noted that 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 config>router>interface>if-attribute
 config>service>ies>interface>if-attribute
 config>service>vprn>interface>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. 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.

It should be noted that 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.

IES Interface IPv6 Commands

ipv6

Syntax	[no] ipv6
Context	config>service>ies>if
Description	This command enables the context to configure IPv6 for an IES interface.

address

Syntax	address <i>ipv6-address/prefix-length</i> [eui-64] no address <i>ipv6-address/prefix-length</i>															
Context	config>service>ies>if>ipv6															
Description	This command assigns an IPv6 address to the IES interface.															
Parameters	<i>ipv6-address/prefix-length</i> — Specify the IPv6 address on the interface.															
Values	<table> <tr> <td>ipv6-address/prefix:</td> <td>ipv6-address</td> <td>x:x:x:x:x:x:x (eight 16-bit pieces)</td> </tr> <tr> <td></td> <td></td> <td>x:x:x:x:x:d.d.d</td> </tr> <tr> <td></td> <td></td> <td>x [0 — FFFF]H</td> </tr> <tr> <td></td> <td></td> <td>d [0 — 255]D</td> </tr> <tr> <td>prefix-length</td> <td></td> <td>1 — 128</td> </tr> </table>	ipv6-address/prefix:	ipv6-address	x:x:x:x:x:x:x (eight 16-bit pieces)			x:x:x:x:x:d.d.d			x [0 — FFFF]H			d [0 — 255]D	prefix-length		1 — 128
ipv6-address/prefix:	ipv6-address	x:x:x:x:x:x:x (eight 16-bit pieces)														
		x:x:x:x:x:d.d.d														
		x [0 — FFFF]H														
		d [0 — 255]D														
prefix-length		1 — 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.															

dad-disable

Syntax	[no] dad-disable
Context	config>service>ies>if>ipv6
Description	<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 no form of the command turns off dad-disable on the interface.</p>
Default	not enabled

icmp6

Syntax	icmp6
Context	config>service>ies>if>ipv6
Description	This command configures ICMPv6 parameters for the IES interface.

packet-too-big

Syntax	packet-too-big [<i>number seconds</i>] no packet-too-big
Context	config>service>ies>if>ipv6>icmp6
Description	This command specifies whether “packet-too-big” ICMPv6 messages should be sent. When enabled, ICMPv6 “packet-too-big” messages are generated by this interface. The no form of the command disables the sending of ICMPv6 “packet-too-big” messages.
Default	100 10
Parameters	<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. Values 10 — 1000 Default 100 <i>seconds</i> — Specifies the time frame in seconds that is used to limit the number of “packet-too-big” ICMPv6 messages issued. Values 1 — 60 Default 10

param-problem

Syntax	param-problem [<i>number seconds</i>] no packet-too-big
Context	config>service>ies>if>ipv6>icmp6
Description	This command specifies whether “parameter-problem” ICMPv6 messages should be sent. When enabled, “parameter-problem” ICMPv6 messages are generated by this interface. The no form of the command disables the sending of “parameter-problem” ICMPv6 messages.
Default	100 10 <i>number</i> — Specifies the number of “parameter-problem” ICMPv6 messages to send in the time frame specified by the <i>seconds</i> parameter.

IES Interface Commands

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of “parameter-problem” ICMPv6 messages issued.

Values 1 — 60

Default 10

redirects

Syntax **redirects** [*number seconds*]
no redirects

Context config>service>ies>if>ipv6>icmp6

Description This command configures ICMPv6 redirect messages. When enabled, ICMPv6 redirects are generated when routes are not optimal on this router and another router on the same subnetwork has a better route in order to alert that node that a better route is available.

When disabled, ICMPv6 redirects are not generated.

Default 100 10

number — Specifies the number of version 6 redirects are to be issued in the time frame specified by the *seconds* parameter.

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of version 6 redirects issued.

Values 1 — 60

Default 10

time-exceeded

Syntax **time-exceeded** [*number seconds*]
no time-exceeded

Context config>service>ies>if>ipv6>icmp6

Description This command specifies whether “time-exceeded” ICMPv6 messages should be sent. When enabled, ICMPv6 “time-exceeded” messages are generated by this interface.

When disabled, ICMPv6 “time-exceeded” messages are not sent.

Default 100 10

number — Specifies the number of “time-exceeded” ICMPv6 messages are to be issued in the time frame specified by the *seconds* parameter.

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of “time-exceeded” ICMPv6 message to be issued.

Values 1 — 60

Default 10

unreachables

Syntax **unreachables** [*number seconds*]
no unreachable

Context config>service>ies>if>ipv6>icmp6

Description This command specifies that ICMPv6 host and network unreachable messages are generated by this interface.

When disabled, ICMPv6 host and network unreachable messages are not sent.

Default 100 10

number — Specifies the number of destination unreachable ICMPv6 messages are issued in the time frame specified by the *seconds* parameter.

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of destination unreachable ICMPv6 messages to be issued.

Values 1 — 60

Default 10

link-local-address

Syntax **link-local-address** *ipv6-address* [**preferred**]
no link-local-address

Context config>router>if>ipv6
config>service>ies>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.

Note that 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.

IES Interface Commands

Parameters **preferred** — Disables duplicated address detection and sets the address to preferred, even if there is a duplicate address.

local-proxy-nd

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

Context config>service>ies>if>ipv6

Description This command enables local proxy neighbor discovery on the interface.
The **no** form of the command disables local proxy neighbor discovery.

proxy-nd-policy

Syntax **proxy-nd-policy** *policy-name* [*policy-name...*(up to 5 max)]
no proxy-nd-policy

Context config>service>ies>if>ipv6

Description This command applies a proxy neighbor discovery policy for the interface.

Parameters *policy-name* — Specifies an existing neighbor discovery policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. The specified policy name(s) must already be defined.

neighbor

Syntax **neighbor** *ipv6-address mac-address*
no neighbor *ipv6-address*

Context config>service>ies>if>ipv6

Description This command configures IPv6-to-MAC address mapping on the IES interface.

Default none

Parameters *ipv6-address* — The IPv6 address of the interface for which to display information.

Values x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:d.d.d.d
x: [0 — FFFF]H
d: [0 — 255]D
prefix-length [1..128]

mac-address — Specifies the 48-bit MAC address for the IPv6-to-MAC address mapping in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

neighbor-limit

Syntax	neighbor-limit <i>limit</i> [log-only] [threshold <i>percent</i>] 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	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 — 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 — 102400

backup

Syntax	[no] backup <i>ip-address</i>
Context	config>service>ies>if>ipv6>vrrp
Description	This command configures virtual router IP addresses for the interface.

init-delay

Syntax	init-delay <i>seconds</i> no init-delay
Context	config>service>ies>if>ipv6>vrrp
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 — 65535

IES Interface Commands

mac

Syntax	mac <i>mac-address</i> no mac
Context	config>service>ies>if>ipv6>vrrp
Description	This command assigns a specific MAC address to an IES IP interface. The no form of the command returns the MAC address of the IP interface to the default value.
Default	The physical MAC address associated with the Ethernet interface that the SAP is configured on (the default MAC address assigned to the interface, assigned by the system).
Parameters	<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

Syntax	[no] master-int-inherit
Context	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 {[<i>seconds</i>] [milliseconds <i>milliseconds</i>]} 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 <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	<i>seconds</i> — The number of seconds that will transpire before the advertisement timer expires. Values 1 — 255 Default 1

milliseconds *milliseconds* — Specifies the time interval, in milliseconds, between sending advertisement messages.

Values 100 — 900

ping-reply

Syntax **[no] ping-reply**

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

IES Interface Commands

Default None

Parameters *vrp-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 [no] preempt

Context config>service>ies>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** *base-priority*
no priority

Context config>service>ies>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 — 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.0

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.

IES Interface Commands

Default no traceroute-reply

IES Spoke SDP Commands

spoke-sdp

Syntax [no] **spoke-sdp** *sdp-id[:vc-id]* [*vc-type* {**ether** | **ipipe**}] [**create**]

Context config>service>ies>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 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.

Default No *sdp-id* is bound to a service.

Restrictions **IES** — At most, only one *sdp-id* can be bound to an IES service.

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 — 4294967295

vc-type — The encapsulation and pseudowire type for the spoke-sdp.

Values ether : Specifies Ethernet pseudowire as the type of virtual circuit (VC) associated with the SDP binding .

ipipe : Specifies Ipipe pseudowire as the type of virtual circuit (VC) associated with the SDP binding .

Default ether

egress

Syntax **egress**

Context config>service>ies>>if>spoke-sdp

Description This command configures the egress SDP context.

QoS

Syntax `qos network-policy-id port-redirect-group queue-group-name [instance instance-id]`
`no qos [network-policy-id]`

Context `configure>service>apipe>spoke-sdp>egress`
`configure>service>cpipe>spoke-sdp>egress`
`configure>service>epipe>spoke-sdp>egress`
`configure>service>fpipe>spoke-sdp>egress`
`configure>service>ipipe>spoke-sdp>egress`
`config>service>vpls>spoke-sdp>egress`
`config>service>vpls>mesh-sdp>egress`
`config>service>pw-template>egress`
`config>service>vprn>interface>spoke-sdp>egress`
`config>service>ies>interface>spoke-sdp>egress`

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-SPD inside a service or to the egress context of a pseudowire template and specify the redirect queue-group name.

One or more spoke-SPDs 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-SPD 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-SPD 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:
 - 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.
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/dot1.p/DSCP and the tunnel DEI/dot1.p/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/dot1.p and the tunnel DEI/dot1.p/EXP, but the DSCP is not modified by the policer operation.

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/dot1.p/DSCP and the tunnel DEI/dot1.p/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 — 65535

queue-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 — 16384

vc-label

Syntax	[no] vc-label egress-vc-label
Context	config>service>ies>if>spoke-sdp>egress
Description	This command configures the static MPLS VC label used by this device to send packets to the far-end device in this service via this SDP.
Parameters	<i>egress-vc-label</i> — A VC egress value that indicates a specific connection.
	Values 16 — 1048575

hash-label

Syntax	hash-label [signal-capability] no hash-label
Context	config>service>ies>if>spoke-sdp
Description	<p>This command enables the use of the hash label on a VLL, VPLS, or VPRN service bound to LDP or RSVP SDP as well as to a VPRN service using the autobind mode with the <i>ldp</i>, <i>rsvp-te</i>, or <i>mpls</i> options. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the <i>gre</i> option.</p> <p>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.</p> <p>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.</p> <p>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. Note however that 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.</p> <p>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.</p> <p>The TTL of the hash label is set to a value of 0.</p> <p>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:</p> <ul style="list-style-type: none">• 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 user 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.

Default no hash-label

Parameters **signal-capability** — Enables the signaling and negotiation of the use of the hash label between the local and remote PE nodes. The **signal-capability** option is not supported on a VPRN spoke-sdp.

ingress

Syntax **ingress**

Context config>service>ies>if>spoke-sdp

Description This command configures the ingress SDP context.

flowspec

Syntax **flowspec**
no flowspec

Context config>service>ies>if>spoke-sdp>ingress

Description This command enables flowspec filtering on an IP interface of the base router. Filtering is based on all of the flowspec routes that have been received and accepted by the base router. Ingress traffic on an IP interface can be filtered by both a user-defined ip filter and flowspec. In this case, the user-defined ip filter entries are evaluated before the flowspec routes and the default action of the user-defined ip filter applies as the very last rule.

The **no** form of the command removes flowspec filtering from an IP interface.

Default No interfaces have flowspec enabled.

qos

Syntax `qos network-policy-id fp-redirect-group queue-group-name instance instance-id]`
`no qos`

Context `configure>service>apipe>spoke-sdp>ingress`
`configure>service>cpipe>spoke-sdp>ingress`
`configure>service>epipe>spoke-sdp>ingress`
`configure>service>fpipe>spoke-sdp>ingress`
`configure>service>ipipe>spoke-sdp>ingress`
`config>service>vpls>spoke-sdp>ingress`
`config>service>vpls>mesh-sdp>ingress`
`config>service>pw-template>ingress`
`config>service>vprn>interface>spoke-sdp>ingress`
`config>service>ies>interface>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:

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).
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.
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.
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.
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:

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.
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-SPD 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.

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.
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.
5. If no network QoS policy is applied to the ingress context of the pseudowire, then 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-SPD, 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 — 65535

fp-redirect-group *queue-group-name* — Specifies the name of the queue group template up to 32 characters in length.

IES Interface Commands

ingress-instance *instance-id* — Specifies the identification of a specific instance of the queue-group.

Values 1 — 16384

vc-label

Syntax [no] **vc-label** *ingress-vc-label*

Context config>service>ies>if>spoke-sdp>ingress

Description This command configures the static MPLS VC label used by the far-end device to send packets to this device in this service via this SDP.

Parameters *ingress-vc-label* — A VC ingress value that indicates a specific connection.

Values 2048 — 18431

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

Context config>service>ies>if>spoke-sdp

Description This command configures an accounting-policy.

Parameters *acct-policy-id* — Specifies an accounting policy ID.

Values 1 — 99

app-profile

Syntax **app-profile** *app-profile-name*
no app-profile

Context config>service>ies>if>spoke-sdp

Description This command configures the application profile name.

Parameters *app-profile-name* — Specifies the application profile name.

bfd-enable

Syntax **bfd-enable**
no bfd-enable

Context config>service>ies>interface>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** *name*
no bfd-template

Context config>service>ies>interface>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 *name* — 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>ies>if>spoke-sdp

Description This command enables or disables statistics collection.

IES SAP Commands

sap

Syntax **sap** *sap-id* [**create**]
no sap *sap-id*

Context config>service>ies>if

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 `access` command.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

Note that you can configure an IES 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.

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 — A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. Attempts to create a second SAP on an IP interface will fail and generate an error; the original SAP will not be affected.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 849 for command syntax.

port-id — Specifies the physical port ID in the *slot/mda/port* format.

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the *slot_number/MDA_number/port_number* format. For example *1/1/1* specifies port 1 on MDA 1 in slot 1.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH, 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.

create — Keyword used to create a SAP instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

aarp

Syntax **aarp** *aarpId* **type** *type*
no aarp

Context config>service>ies>if>sap
config>service>ies>if>spoke-sdp

Description This command associates an aarp instance to a multi-homed SAP or spoke-sdp. This instance is paired with the same aarp-id in the same node or 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 specifies the role of this service point in the AARP: primary (dual-homed), secondary (dual-homed-secondary). The AA service attributes (app-profile, 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.

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 —

type — Specifies the role of the SAP referenced by the AARP instance identified by AARP ID.

Values **dual-homed** — the primary dual homed aa-subscriber side service point of an aarp instance, only supported for 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 IES and VPRN SAP and spoke-sdp.

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.

IES Interface Commands

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, etc.), the entire string must be enclosed within double quotes.

lag-link-map-profile

Syntax **lag-link-map-profile** *lag-link-map-profile-id*
no lag-link-map-profile

Context config>service>ies>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 *lag-link-map-profile-id* — 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

Syntax **lag-per-link-hash class {1 | 2 | 3} weight [1..1024]**
no 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)

multi-service-site

Syntax **multi-service-site** *customer-site-name*
no multi-service-site *customer-site-name*

Context config>service>ies>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 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.

Default None — Each customer site must be explicitly created.

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 queues on the SAPs to no longer be orphaned. Existing schedulers on the site may cease to exist, causing 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 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, etc.), the entire string must be enclosed within double quotes.

static-host

Syntax **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*

Context config>service>ies>if>sap

Description This command configures a static host on this SAP.

Parameters **ip** *ip-address* — Specifies the IPv4 unicast address.

IES Interface Commands

mac *ieee-address* — Specify 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.

sla-profile *sla-profile-name* — 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>subscriber>sla-profile** context.

ancp-string

Syntax **ancp-string** *ancp-string*
 no ancp-string

Context config>service>ies>if>sap>static-host

Description This command specifies the ANCP string associated to this SAP host.

Parameters *ancp-string* — Specifies the ANCP string up to 63 characters in length.

app-profile

Syntax **app-profile** *app-profile-name*
 no app-profile

Context config>service>ies>if>sap>static-host

Description This command specifies an application profile name.

Parameters *app-profile-name* — Specifies the application profile name up to 32 characters in length.

inter-dest-id

Syntax **inter-dest-id** *intermediate-destination-id*
 no inter-dest-id

Context config>service>ies>if>sap>static-host

Description Specifies to which intermediate destination (for example, a DSLAM) this host belongs.

Parameters *intermediate-destination-id* — Specifies the intermediate destination identifier, up to 32 characters in length.

sla-profile

Syntax **sla-profile** *sla-profile-name*
 no sla-profile

Context config>service>ies>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** *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
- 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
- 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
- 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.

tod-suite

- Syntax** **tod-suite** *tod-suite-name*
no tod-suite
- Context** config>service>ies>if>sap
- Description** This command applies a time-based policy (filter or QoS policy) to the service SAP. The suite name must already exist in the **config>system>cron** context.

IES Interface Commands

Default no tod-suite

Parameters *tod-suite-name* — Specifies collection of policies (ACLs, QoS) including time-ranges that define the full or partial behavior of a SAP. The suite can be applied to more than one SAP.

transit-policy

Syntax **transit-policy** *ip-aasub-policy-id*
no transit-ip-policy

Context config>service>ies>if>sap>
config>service>ies>if>spoke-sdp>

Description This command associates a transit aa subscriber IP 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 the use of this policy assigned to services, 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-ip-policy

ip-aasub-policy-id — An integer that identifies a transit IP profile entry.

Values 1 — 65535

Values

dynamic-tunnel-redundant-next-hop

Syntax **dynamic-tunnel-redundant-next-hop** *ip-address*
no dynamic-tunnel-redundant-next-hop

Context config>service>ies>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.

Default none

Parameters *ip-address* — 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>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 l4-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 l4-load balancing is enabled) source port in the hash, ignore destination address/port. destination — Specifies using destination address and (if l4-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>interface>sap>ingress config>service>vprn>interface>spoke-sdp>ingress config>service>ies>interface>sap>ingress config>service>ies>interface>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: <ol style="list-style-type: none"> 1. user-defined IPv4 filter entries 2. flowspec-derived filter entries 3. user-defined IPv4 filter default-action <p>The no form of the command removes IPv4 flowspec filtering from an IP interface.</p>
Default	No access interfaces have IPv4 flowspec enabled.

flowspec-ipv6

Syntax	flowspec-ipv6 no flowspec-ipv6
Context	config>service>vprn>interface>sap>ingress config>service>vprn>interface>spoke-sdp>ingress config>service>ies>interface>sap>ingress config>service>ies>interface>spoke-sdp>ingress
Description	This command enables IPv6 flowspec filtering on an access IP interface associated with a VPRN or IES service. Filtering is based on all of the IPv6 flowspec routes that have been received and accepted by the corresponding BGP instance. Ingress IPv6 traffic on an interface can be filtered by both a user-defined IPv6 filter and flowspec. Evaluation proceeds in this order: <ol style="list-style-type: none"> 1. user-defined IPv6 filter entries 2. flowspec-derived filter entries 3. user-defined IPv6 filter default-action <p>The no form of the command removes IPv6 flowspec filtering from an IP interface.</p>
Default	No access interfaces have IPv6 flowspec enabled.

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	config>service>ies>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. Note that 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 <i>interval</i> — 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 — 6000 Note that 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 etc.). 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— 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— 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— 29

SOURCE

Syntax **source** *ip-address*

Context config>service>ies>interface>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 *ip-address* — An IPv4 address or an IPv6 address.

remote-ip

Syntax **remote-ip** *ip-address*
no remote-ip

Context config>service>ies>interface>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 *ip-address* — 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>interface>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 remote-ip
Parameters	<i>ip-address</i> — An IPv4 address or an IPv6 address.

accounting-policy

Syntax	accounting-policy <i>acct-policy-id</i> no accounting-policy
Context	config>service>ies>if>sap
Description	This command creates the accounting policy context that can be applied to a SAP. An accounting policy must be defined before it can be associated with a SAP. If the <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 context. The no form of this command removes the accounting policy association from the SAP, and the accounting policy reverts to the default.
Default	Default accounting policy.
Parameters	<i>acct-policy-id</i> — Enter the accounting <i>policy-id</i> as configured in the config>log>accounting-policy context. Values 1 to 99

calling-station-id

Syntax	calling-station-id <i>calling-station-id</i> no calling-station-id
Context	config>service>ies>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 This attribute is not sent by default.

ETH-CFM Service Commands

eth-cfm

Syntax	eth-cfm
Context	config>service>ies> config>service>ies>if>sap config>service>ies>if>spoke-sdp
Description	This command enables the context to configure ETH-CFM parameters.

collect-lmm-stats

Syntax	collect-lmm-stats no collect-lmm-stats
Context	config>service>ies>interface>sap>>eth-cfm config>service>ies>interface>spoke-sdp>>eth-cfm config>service>ies>subscriber-interface>group-interface>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

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
Description	This command configures the ETH-CFM maintenance endpoint (MEP).
Parameters	<i>mep-id</i> — Specifies the maintenance association end point identifier. Values 1 — 8191 <i>md-index</i> — Specifies the maintenance domain (MD) index value.

Values 1 — 4294967295

ma-index — Specifies the MA index value.

Values 1 — 4294967295

direction up|down — Indicates 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
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep

Description This command configures the reception of Alarm Indication Signal (AIS) message.

interface-support-enable

Syntax [no] **interface-support-enable**

Context config>service>ies>sap>eth-cfm>mep>ais-enable
config>service>ies>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.

Default [no] interface-support-enabled (AIS will not be generated or stopped based on the state of the entity on which the DOWN MEP is configured).

ccm-enable

Syntax [no] **ccm-enable**

Context config>service>ies>if>sap>eth-cfm>mep
config>service>ies>if>spoke-sdp>eth-cfm>mep

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 <i>priority</i> no ccm-ltm-priority
Context	config>service>ies>if>sap>eth-cfm>mep config>service>ies>if>spoke-sdp>eth-cfm>mep
Description	This command specifies the priority value for CCMs and LTMs transmitted by the MEP. The no form of the command removes the priority value from the configuration.
Default	The highest priority on the bridge-port.
Parameters	<i>priority</i> — Specifies the priority of CCM and LTM messages. Values 0 — 7

ccm-padding-size

Syntax	[no] ccm-padding-size <i>ccm-padding</i>
Context	config>service>ies>if>spoke-sdp>eth-cfm>mep
Description	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.
Default	ccm-padding-size
Parameters	<i>ccm-padding</i> — specifies the byte size of the Optional Data TLV Values 3 — 1500

eth-test-enable

Syntax	[no] eth-test-enable
Context	config>service>ies>if>sap>eth-cfm>mep
Default	For ETH-test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands: oam eth-cfm eth-test <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> 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>ies>if>sap>eth-cfm>mep>eth-test-enable
Default	This command configures the test pattern for eth-test frames. The no form of the command removes the values from the configuration.
Parameters	all-zeros — Specifies to use all zeros in the test pattern. all-ones — Specifies to use all ones in the test pattern. crc-enable — Generates a CRC checksum.
	Default all-zeros

fault-propagation-enable

Syntax	fault-propagation-enable {use-if-tlv suspend-ccm} no fault-propagation-enable
Context	config>service>ies>if>sap>eth-cfm>mep config>service>ies>if>spoke-sdp>eth-cfm>mep
Description	This command configures the fault propagation for the MEP.
Parameters	use-if-tlv — Specifies to use the interface TLV. suspend-ccm — Specifies to suspend the continuity check messages.

low-priority-defect

Syntax	low-priority-defect {allDef macRemErrXcon remErrXcon errXcon xcon noXcon}
Context	config>service>ies>if>sap>eth-cfm>mep
Description	This command specifies the lowest priority defect that is allowed to generate a fault alarm.
Default	macRemErrXcon
	Values
	allDef DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
	macRemErrXcon Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
	remErrXcon Only DefRemoteCCM, DefErrorCCM, and DefXconCCM
	errXcon Only DefErrorCCM and DefXconCCM
	xcon Only DefXconCCM; or
	noXcon No defects DefXcon or lower are to be reported

squelch-ingress-levels

Syntax	squelch-ingress-levels [<i>md-level</i> [<i>md-level</i> ...]] no squelch-ingress-levels
Context	config>service>ies>interface>sap>eth-cfm config>service>ies>interface>spoke-sdp>eth-cfm config>service>ies>subscriber-interface>group-interface>sap>eth-cfm
Description	<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 no form of the command removes the silent discarding of previously matching ETH-CFM PDUs.</p>
Default	no squelch-ingress-levels
Parameters	<i>md-level</i> — Identifies the level.
	Values [0..7]

tunnel-fault

Syntax	tunnel-fault { accept ignore }
Context	config>service>ies>eth-cfm config>service>ies>if>sap>eth-cfm
Description	<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 command ais-enable under epipe>sap>eth-cfm>ais-enable for more details. 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.</p>
Parameters	<p>accept — Share fate with the facility tunnel MEP</p> <p>ignore — Do not share fate with the facility tunnel MEP</p>
Default	<p>ignore (Service Level)</p> <p>accept (SAP Level for Epipe and VPLS)</p>

one-way-delay-threshold

Syntax `one-way-delay-threshold time`

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

Description This command enables one way delay threshold time limit.

Default 3 seconds

Parameters *priority* — Specifies the value for the threshold.

Values 0 — 600

IES Filter and QoS Policy Commands

filter

Syntax

```
filter ip ip-filter-id
filter ipv6 ipv6-filter-id
no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
```

Context

```
config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress
config>service>ies>redundant-interface>egress
config>service>ies>redundant-interface>ingress
```

Description

This command associates a filter policy with an ingress or egress Service Access Point (SAP). Filter policies control the forwarding and dropping of packets based on the matching criteria.

The **filter** command is used to associate a filter policy with a specified *ip-filter-id* with an ingress or egress SAP. The filter policy must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to the match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP. The filter ID 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**.

Restrictions

IES — Only IP filters are supported on an IES IP interface, and the filters only apply to routed traffic.

Parameters

ip — Keyword indicating the filter policy is an IP filter.

ip-filter-id — Specifies the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy in the **configure>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 **configure>filter** context before the **filter** command is executed. If the filter policy does not exist, the operation will fail 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**.

- Specifies** **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

Syntax **egress**

Context config>service>ies>if>sap

Description This command enables the context to apply egress 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>ies>if>sap

Description This command enables the context to apply ingress 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.

match-qinq-dot1p

Syntax `match-qinq-dot1p {top | bottom}`
`no match-qinq-dot1p`

Context `config>service>ies>if>sap>ingress`

Description This command specifies which Dot1Q tag position 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 4](#) defines the default behavior for Dot1P evaluation when the **match-qinq-dot1p** command is not executed.

Table 4: 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
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 5](#) defines the dot1p evaluation behavior when the top parameter is specified.

Table 5: 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 6](#) defines the dot1p evaluation behavior when the bottom parameter is specified.

Table 6: 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

Table 6: Bottom Position QinQ and TopQ SAP Dot1P Evaluation (Continued)

Port / SAP Type	Existing Packet Tags	PBits Used for Match
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	BottomQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

Table 7: 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 8: 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

Table 8: QinQ Mark Top Only Explicit Marking Actions

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
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 the above tables 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-limit

Syntax **agg-rate-limit** *agg-rate* [**queue-frame-based-accounting**]
no agg-rate-limit

Context config>service>ies>if>sap>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 will fail. If the **agg-rate-limit** command is specified, an attempt to bind a **scheduler-policy** to the SAP or multi-service site will fail.

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 will fail. 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 are 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 MSS can operate.

Values 1 — 40000000, max

queue-frame-based-accounting — This keyword enables frame based accounting on all queues associated

IES Filter and QoS Policy Commands

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

agg-rate

Syntax [no] **agg-rate**

Context config>service>ies>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

Syntax **rate** {max | rate}
no rate

Context config>service>ies>if>sap>egress>agg-rate

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 etc.).

limit-unused-bandwidth

Syntax [no] **limit-unused-bandwidth**

Context config>service>ies>if>sap>egress>agg-rate

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>ies>if>sap>egress>agg-rate

Description This command is used to enabled (or disable) frame based accounting on all queues associated with the agg-rate context. Only supported on Ethernet ports. Not supported on HSMDA Ethernet ports.

qinq-mark-top-only

Syntax [no] qinq-mark-top-only

Context config>service>ies>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 *policy-id* [**port-redirect-group** *queue-group-name* **instance** *instance-id*]
no qos

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

Description This command associates a Quality of Service (QoS) policy with an 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. 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 at one time. Attempts to associate a second QoS policy of a given type will return an error.

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.

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.

Default none

Parameters *policy-id* — The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.

1 — 65535

port-redirect-group — 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.

queue-group-name — Specifies the name of the egress port queue group of the IOM/IMM/XMA, up to 32

IES Filter and QoS Policy Commands

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 — 40960

Default 1

qos

Syntax **qos** *policy-id* [**shared-queuing** | **multipoint-shared**] [**fp-redirect-group** *queue-group-name* **instance** *instance-id*]
no qos

Context config>service>vprn>if>sap>ingress
config>service>vprn>sub-if>grp-if>sap>ingress
config>service>vprn>ipsec-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.

The **no** form of this command removes the QoS policy association from the SAP or IP interface, and the QoS policy reverts to the default.

Default none

Parameters *policy-id* — The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.

1 — 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 — 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.

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 — This keyword 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-id are mandatory parameters when executing the command. The named queue group template can contain only policers. If it contains queues, then the command will fail.

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.

queue-override

Syntax [no] queue-override

Context config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress

Description This command enables 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.

queue

Syntax [no] queue *queue-id*

Context config>service>ies>if>sap>egress>queue-override
config>service>ies>if>sap>ingress>queue-override

Description This command specifies the ID of the queue whose parameters are to be overridden.

Parameters *queue-id* — The queue ID whose parameters are to be overridden.

Values 1 — 32

adaptation-rule

Syntax	adaptation-rule [pir { max min closest }] [cir { max min closest }] no adaptation-rule
Context	config>service>ies>if>sap>egress>queue-override>queue config>service>ies>if>sap>ingress>queue-override>queue
Description	<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>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.</p>
Default	no adaptation-rule
Parameters	<p>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.</p> <p>max — The max (maximum) option 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.</p> <p>min — The min (minimum) option 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.</p> <p>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.</p> <p>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.</p>

avg-frame-overhead

Syntax	avg-frame-overhead <i>percent</i> no avg-frame-overhead
Context	config>service>ies>if>sap>egress>queue-override config>service>ies>if>sap>ingress>queue-override>queue
Description	<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</p>

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 7500×1.1 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.

IES Filter and QoS Policy Commands

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 — 100

cbs

Syntax **cbs** *size-in-kbytes*
no cbs

Context config>service>ies>if>sap>egress>queue-override>queue
config>service>ies>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 size to the default value.

Default no cbs

Parameters *size-in-kbytes* — The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

Values 0 — 131072, default

high-prio-only

Syntax **high-prio-only** *percent*
no high-prio-only

Context config>service>ies>if>sap>egress>queue-override>queue
config>service>ies>if>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's high-prio-only parameters. The **high-prio-only** command configures the percentage of buffer space for the queue, used exclusively by high priority packets.

The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The **high-prio-only** parameter is used to override the default value derived from the **network-queue** command.

The defined **high-prio-only** value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the **high-prio-only** value larger than the current MBS size will also result in an error and fail execution.

The **no** form of this command restores the default high priority reserved size.

Parameters *percent* — The *percent* parameter is the percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that none of the MBS of the queue will be reserved for high priority traffic. This does not affect RED slope operation for packets attempting to be queued.

Values 0 — 100, default

mbs

Syntax **mbs** {*size-in-kbytes* | **default**}
no mbs

Context config>service>ies>if>sap>egress>queue-override>queue
config>service>ies>if>sap>egress>hsmda-queue-override>queue
config>service>ies>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 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

IES Filter and QoS Policy Commands

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 default

Parameters *size-in-kbytes* — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.

Values 0 — 131072 or default

monitor-depth

Syntax **monitor-depth**
no monitor-depth

Context config>service>ies>if>sap>egress>queue-override>queue
config>service>ies>if>sap>ingress>queue-override>queue
config>service>vprn>if>sap>egress>queue-override>queue
config>service>vprn>if>sap>ingress>queue-override>queue

Description This command enables queue depth monitoring for the specified queue.
The **no** form of the command removes queue depth monitoring for the specified queue.

parent

Syntax **parent [weight weight] [cir-weight cir-weight]**
no parent

Context config>service>ies>if>sap>ingress>sched-override>scheduler
config>service>ies>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	<p>weight <i>weight</i> — 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.</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>Values 0 to 100</p> <p>Default 1</p> <p>cir-weight <i>cir-weight</i> — 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 queue or scheduler that the cir-weight pertains to has not reached the CIR and still has packets to transmit.</p> <p>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.</p> <p>Values 0 — 100</p> <p>Default 1</p>
-------------------	---

rate

Syntax	<p>rate <i>pir-rate</i> [<i>cir cir-rate</i>]</p> <p>no rate</p>
Context	<p>config>service>ies>if>sap>egress>queue-override>queue</p> <p>config>service>ies>if>sap>ingress>queue-override>queue</p> <p>config>service>ies>if>sap>egress>sched-override>scheduler</p> <p>config>service>ies>if>sap>ingress>sched-override>scheduler</p>
Description	<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 packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile 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 rate command can be executed at any time, altering the PIR and CIR rates for all queues created</p>

IES Filter and QoS Policy Commands

through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default **rate max cir 0** — The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir-rate** value.

Parameters *pir-rate* — 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.

For **egress>queue-override>queue** and **ingress>queue-override>queue**:

Values 1 — 2000000000, **max** in Kbps

Default max

For **egress>sched-override>scheduler** and **ingress>sched-override>scheduler**:

Values 1 — 3200000000, **max** in Kbps

Default max

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— 2000000000, **sum**, **max** in Kbps

Default 0

For **egress>sched-override>scheduler** and **ingress>sched-override>scheduler**:

Values 0— 3200000000, **sum**, **max** in Kbps

Default 0

scheduler-override

Syntax [**no**] **scheduler-override**

Context config>service>ies>if>sap>egress
config>service>ies>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 [no] scheduler *scheduler-name*

Context config>service>ies>if>sap>egress>sched-override
config>service>ies>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 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 queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier, except for schedulers created in tier 1). 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:

1. The maximum number of schedulers has not been configured.
2. The provided *scheduler-name* is valid.
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, etc.), 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 all queues created with this *queue-id* by association with the QoS policy to the default PIR and CIR parameters.

Parameters *pir-rate* — The **pir** parameter accepts a step multiplier value that specifies the multiplier used to determine the PIR rate at which the queue will operate. A value of 0 to 100000000 or the keyword **max** is accepted. Any other value will result in an error without modifying the current PIR rate.

To calculate the actual PIR rate, the rate described by the queue’s **rate** is multiplied by the *pir-rate*.

The SAP ingress context for PIR is independent of the defined forwarding class (fc) for the queue. The default **pir** and definable range is identical for each class. The PIR in effect for a queue defines the maximum rate at which the queue will be allowed to forward packets in a given second, thus shaping the queue’s output.

The PIR parameter for SAP ingress queues do not have a negate (**no**) function. To return the queue’s PIR rate to the default value, that value must be specified as the PIR value.

Values 1 — 100000000, **max**

Default 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 0 — 100000000 or the keyword **max** or **sum** is accepted. Any other value will result in an error without modifying the current CIR rate.

To calculate the actual CIR rate, the rate described by the **rate** **pir** *pir-rate* is multiplied by the *cir* *cir-rate*. If the **cir** is set to max, then the CIR rate is set to infinity.

The SAP ingress context for CIR is dependent on the defined forwarding class (fc) for the queue. The default CIR and definable range is different for each class. The CIR in effect for a queue defines both its profile (in or out) marking level as well as the relative importance compared to other queues for scheduling purposes during congestion periods. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues.

For **egress>sched-override>scheduler** and **ingress>sched-override>scheduler**:

Values 0 — 2000000000, **max**, **sum**

Default sum

scheduler-policy

Syntax **scheduler-policy** *scheduler-policy-name*
no scheduler-policy

Context config>service>ies>sap>ingress
config>service>ies>sap>egress

Description 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 **config>qos>scheduler-policy** *scheduler-policy-name* 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 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 **no scheduler-policy** command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.

scheduler-policy-name: — The *scheduler-policy-name* 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.

IES Interface VRRP Commands

vrrp

Syntax	vrrp <i>virtual-router-id</i> [owner] no vrrp <i>virtual-router-id</i>
Context	config>service>ies>if
Description	<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 owner. The nodal context of vrrp <i>virtual-router-id</i> is used to define the configuration parameters for the VRID.</p> <p>The no form of this command removes the specified VRID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the vrid. The VRID does not need to be shutdown in order to remove the virtual router instance.</p>
Default	No default
Parameters	<p><i>virtual-router-id</i> — The virtual-router-id parameter specifies a new virtual router ID or one that can be modified on the IP interface.</p> <p>Values 1 — 255</p>

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>service>ies>if>vrrp
Description	<p>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 validating received VRRP advertisement messages.</p> <p>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.</p> <p>To change the current in-use password key on multiple virtual router instances:</p> <ul style="list-style-type: none"> • 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** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

backup

Syntax [no] **backup** *ip-address*

Context config>service>ies>if>vrrp

Description This command configures virtual router IP addresses for the interface.

bfd-enable

Syntax [no] **bfd-enable** [*service-id*] **interface** *interface-name* **dst-ip** *ip-address*

Context config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

IES Filter and QoS Policy Commands

Description	<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 no form of this command removes BFD from the configuration.</p>
Default	none
Parameters	<p><i>service-id</i> — Specifies the service ID of the interface running BFD.</p> <p>Values <i>service-id</i>: 1 — 2147483648</p> <p>Values No service ID indicates a network interface.</p> <p>interface <i>interface-name</i> — Specifies the name of the interface running BFD.</p> <p>dst-ip <i>ip-address</i> — Specifies the destination address to be used for the BFD session.</p>

init-delay

Syntax	init-delay <i>seconds</i> no init-delay
Context	config>service>ies>if>vrrp
Description	This command configures a VRRP initialization delay timer.
Default	no init-delay
Parameters	<p><i>seconds</i> — Specifies the initialization delay timer for VRRP, in seconds.</p> <p>Values 1 — 65535</p>

mac

Syntax	mac <i>mac-address</i> no mac
Context	config>service>ies>if>vrrp
Description	<p>This command assigns a specific MAC address to an IES IP interface.</p> <p>The no form of the command returns the MAC address of the IP interface to the default value.</p>
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	<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>

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

Syntax	message-interval {[seconds] [milliseconds milliseconds]} no message-interval
Context	config>service>ies>if>vrrp
Description	<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 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.</p> <p>The no form of this command restores the default message interval value of 1 second to the virtual router instance.</p>
Parameters	<p><i>seconds</i> — The number of seconds that will transpire before the advertisement timer expires.</p> <p>Values 1 — 255</p> <p>Default 1</p> <p><i>milliseconds milliseconds</i> — Specifies the time interval, in milliseconds, between sending advertisement messages. This parameter is not supported on non-redundant chassis.</p> <p>Values 100 — 900</p>

ping-reply

Syntax	ping-reply no ping-reply
Context	config>service>ies>if>vrrp
Description	<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>

IES Filter and QoS Policy Commands

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 None

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 affect of the dynamic changing of the in-use priority is greatly diminished.

The preempt command is only available in the non-owner vrrp virtual-router-id nodal context. The owner may not be preempted due to the fact that the priority of non-owners can never be higher than the owner. The owner 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

Syntax **priority** *base-priority*
no priority

Context config>service>ies>if>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 — 254

Default 100

standby-forwarding

Syntax [**no**] **standby-forwarding**

Context config>service>ies>if>vrrp

Description This command allows the forwarding of packets by a standby router.

IES Filter and QoS Policy Commands

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

ssh-reply

Syntax [no] ssh-reply

Context config>service>ies>if>vrrp

Description This command enables the non-owner master to reply to SSH Requests directed at the 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

Syntax [no] traceroute-reply

Context config>service>ies>if>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

IES Show Commands

customer

Syntax `customer [customer-id] [site customer-site-name]`

Context `show>service`

Description This command displays service customer information.

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

Default All customer IDs display

Values 1 — 2147483647

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

Output **Show Customer Command Output** — The following table describes show customer command output fields:

Label	Description
Customer-ID	The ID that uniquely identifies a customer.
Contact	The name of the primary contact person.
Description	Generic information about the customer.
Phone	The phone/pager number to reach the primary contact person.
Total Customers	The total number of customers configured.
Multi-service site	
Site	Multi-service site name. A multi-service customer site is a group of SAPs with common origination and termination points.
Description	Information about a specific customer's multi-service site.
Assignment	The port ID, MDA, or card number, where the 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	Specifies the SAP assigned to the service.

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#

```

sap-using

Syntax

```

sap-using [msap] [dyn-script] [description]
sap-using [sap sap-id] [vlan-translation | anti-spoof] [description]
sap-using [sap sap-id]
sap-using interface [ip-address | ip-int-name]
sap-using [ingress | egress] filter filter-id
sap-using [ingress ] qos-policy qos-policy-id

```

Context show>service

Description Displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs. The optional parameters restrict output to only SAPs matching the specified properties.

Parameters **sap** *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 849 for command syntax.

ingress — Specifies matching an ingress policy.

egress — Specifies matching an egress policy.

filter *filter-id* — The ingress or egress filter policy ID for which to display matching SAPs.

Values 1 — 65535

IES Show Commands

dyn-script — Displays dynamic service SAPs information.

interface — Specifies matching SAPs with the specified IP interface.

ip-addr — The IP address of the interface for which to display matching SAPs.

Values 1.0.0.0 — 223.255.255.255

ip-int-name — The IP interface name for which to display matching SAPs.

Output **Show Service SAP** — The following table describes show service SAP output fields:

Label	Description
Port ID	The ID of the access port where the SAP is defined.
Svc ID	The value that identifies the service.
Ingr.QoS	The SAP ingress QoS policy number specified on the ingress SAP.
Ing.Fltr	The MAC or IP filter policy ID applied to the ingress SAP.
Egr.Fltr	The MAC or IP filter policy ID applied to the egress SAP.
Adm	The administrative state of the SAP.
Opr	The actual state of the SAP.

Sample Output

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

service-using

Syntax `service-using [ies] [customer customer-id]`

Context `show>service`

Description This command displays the services matching certain usage properties. If no optional parameters are specified, all services defined on the system are displayed.

Parameters `ies` — Displays matching IES services.

`customer customer-id` — Displays services only associated with the specified customer ID.

Default Services associated with an customer.

Values 1 — 2147483647

Output **Show Service Service-Using** — The following table describes show service service-using output fields:

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.

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#
```

id

Syntax `id service-id {all | arp | base | sap}`

Context show>service

Description This command displays information for a particular service-id.

Parameters *service-id* — The unique service identification number to identify the service in the service domain.

all — Display detailed information about the service.

arp — Display ARP entries for the service.

base — Display basic service information.

interface — Display service interfaces.

sap — Display SAPs associated to the service.

all

Syntax `all`

Context show>service>id

Description This command displays detailed information for all aspects of the service.

Output **Show All Service-ID Output** — The following table describes the show all service-id command output fields:

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.

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.

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: “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.

Label	Description (Continued)
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.
Queuing 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.

Label	Description (Continued)
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)

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
tod-suite      : None                  qinq-pbit-marking : both
Egr Agg Rate Limit: max

Endpoint       : N/A
Limit Unused BW : Disabled

Vlan-translation : None
Acct. Pol       : None                 Collect Stats   : Disabled
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

-----
...

*A:Dut-B# show service id 1 all

=====
Service Detailed Information
=====
Service Id      : 1                    Vpn Id         : 0
Service Type    : IES
```

```

Name          : (Not Specified)
Description   : (Not Specified)
Customer Id   : 1                Creation Origin  : manual
Last Status Change: 01/28/2015 22:17:56
Last Mgmt Change  : 01/28/2015 22:10:04
Admin State    : Up              Oper State      : Up
SAP Count      : 0              SDP Bind Count  : 1

```

 ETH-CFM service specifics

Tunnel Faults : ignore

 Service Destination Points(SDPs)

Sdp Id 230:1 -(10.20.1.3)

```

Description   : (Not Specified)
SDP Id        : 230:1                Type           : Spoke
Spoke Descr   : (Not Specified)
VC Type       : Ether                VC Tag         : n/a
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
Hash Label    : Disabled            Hash Lbl Sig Cap : Disabled
Oper Hash Label : Disabled

```

```

Admin State    : Up              Oper State      : Up
Acct. Pol     : None            Collect Stats   : Disabled
Ingress Label  : 262133         Egress Label   : 262133
Ingr Mac Fltr-Id : n/a          Egr Mac Fltr-Id : n/a
Ingr IP Fltr-Id : n/a          Egr IP Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a        Egr IPv6 Fltr-Id : n/a
BGP IPv4 FlowSpec : Disabled
BGP IPv6 FlowSpec : Disabled
Admin ControlWord : Not Preferred Oper ControlWord : False
BFD Template   : None
BFD-Enabled    : no            BFD-Encap      : ipv4
Last Status Change : 01/28/2015 22:17:56 Signaling       : TLDP
Last Mgmt Change  : 01/28/2015 22:16:50
Class Fwding State : Down
Flags          : None
Local Pw Bits   : None
Peer Pw Bits    : None
Peer Fault Ip   : None
Peer Vccv CV Bits : lspPing bfdFaultDet
Peer Vccv CC Bits : mplsRouterAlertLabel

```

```

Application Profile: None
Transit Policy     : None
AARP Id           : None

```

```

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

```

KeepAlive Information :

IES Show Commands

```
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.   : 3                E. Fwd. Octets  : 180
```

Control Channel Status

```
PW Status       : disabled          Refresh Timer    : <none>
Peer Status Expire : false
Request Timer    : <none>
Acknowledgement  : false
```

ETH-CFM SDP-Bind specifics

```
Squelch Levels  : 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
```

```
Number of SDPs : 1
```

Service Access Points

```
No Sap Associations
```

Service Interfaces

Interface

```
If Name          : iesSpokeToC
Admin State      : Up                Oper (v4/v6)     : Up/Down
Protocols        : None
IP Addr/mask     : 20.20.20.2/24      Address Type     : Primary
IGP Inhibit      : Disabled          Broadcast Address : Host-ones
```

HoldUp-Time : 0 Track Srrp Inst : 0
Description : N/A

Details

Description : (Not Specified)
If Index : 4 Virt. If Index : 4
Last Oper Chg : 01/28/2015 22:17:56 Global If Index : 257
Mon Oper Grp : None
Srrp En Rtng : Disabled Hold time : N/A
SDP Id : spoke-230:1

Spoke-SDP Details

Admin State : Up Oper State : Up
Hash Label : Disabled Hash Lbl Sig Cap : Disabled
Oper Hash Label : Disabled
Peer Fault Ip : None
Local Pw Bits : None
Peer Pw Bits : None
Peer Vccv CV Bits : lspPing bfdFaultDet
Peer Vccv CC Bits : mplsRouterAlertLabel
Flags : None

TOS Marking : Untrusted If Type : IES
SNTP B.Cast : False IES ID : 1
MAC Address : 00:03:fa:32:16:62 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
IP Oper MTU : 1500 ICMP Mask Reply : True
Arp Populate : Disabled Host Conn Verify : Disabled
Cflowd (unicast) : None Cflowd (multicast): None
LdpSyncTimer : None
LSR Load Balance : system
EGR Load Balance : both
TEID Load Balance : Disabled
SPI Load Balance : Disabled
uRPF Chk : disabled
uRPF Ipv6 Chk : disabled
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 Pkts : 3 Tx Bytes : 180
Tx V4 Pkts : 0 Tx V4 Bytes : 0
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

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

IES Show Commands

Description : (Not Specified)
Admin State : Down Lease Populate : 0
Gi-Addr : 20.20.20.2* 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

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

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:Dut-B#

*B:Dut-C# show service id 1 all

=====
Service Detailed Information
=====

Service Id : 1 Vpn Id : 1
Service Type : IES
Name : XYZ Ies 1
Description : default
Customer Id : 1 Creation Origin : manual
Last Status Change: 05/27/2015 07:15:06
Last Mgmt Change : 05/27/2015 07:14:43
Admin State : Up Oper State : Up
SAP Count : 0 SDP Bind Count : 1

ETH-CFM service specifics

Tunnel Faults : ignore

Service Destination Points(SDPs)

Sdp Id 6:1 - (10.20.1.1)

Description : Default sdp description
SDP Id : 6:1 Type : Spoke
Spoke Descr : (Not Specified)
VC Type : Ether VC Tag : n/a
Admin Path MTU : 1514 Oper Path MTU : 1514
Delivery : MPLS
Far End : 10.20.1.1
Tunnel Far End : n/a LSP Types : SR-OSPF
Hash Label : Disabled Hash Lbl Sig Cap : Disabled
Oper Hash Label : Disabled

Admin State : Up Oper State : Up
Acct. Pol : None Collect Stats : Disabled
Ingress Label : 262133 Egress Label : 262134
Ingr Mac Fltr-Id : n/a Egr Mac Fltr-Id : n/a
Ingr IP Fltr-Id : n/a Egr IP Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a Egr IPv6 Fltr-Id : n/a
BGP IPv4 FlowSpec : Disabled
BGP IPv6 FlowSpec : Disabled
Admin ControlWord : Not Preferred Oper ControlWord : False
BFD Template : None
BFD-Enabled : no BFD-Encap : ipv4
Last Status Change : 05/27/2015 07:15:06 Signaling : TLDP
Last Mgmt Change : 05/27/2015 07:14:43
Class Fwding State : Down
Flags : None
Local Pw Bits : None
Peer Pw Bits : None
Peer Fault Ip : None
Peer Vccv CV Bits : lspPing bfdFaultDet
Peer Vccv CC Bits : mplsRouterAlertLabel

No Sap Associations

Service Interfaces

Interface

If Name : ies-1
Admin State : Up Oper (v4/v6) : Up/Down
Protocols : None
IP Addr/mask : 10.3.1.3/24 Address Type : Primary
IGP Inhibit : Disabled Broadcast Address : Host-ones
HoldUp-Time : 0 Track Srrp Inst : 0
Description : N/A

Details

Description : (Not Specified)
If Index : 3 Virt. If Index : 3
Last Oper Chg : 05/27/2015 07:15:06 Global If Index : 257
Mon Oper Grp : None
Srrp En Rtng : Disabled Hold time : N/A
SDP Id : spoke-6:1

Spoke-SDP Details

Admin State : Up Oper State : Up
Hash Label : Disabled Hash Lbl Sig Cap : Disabled
Oper Hash Label : Disabled
Peer Fault Ip : None
Local Pw Bits : None
Peer Pw Bits : None
Peer Vccv CV Bits : lspPing bfdFaultDet
Peer Vccv CC Bits : mplsRouterAlertLabel
Flags : None

TOS Marking : Untrusted If Type : IES
SNTP B.Cast : False IES ID : 1
MAC Address : 66:34: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
IP MTU : 1400
IP Oper MTU : 1400 ICMP Mask Reply : True
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

IES Show Commands

```
PTP HW Assist      : Disabled
Rx Pkts           : 100
Rx V4 Pkts        : N/A
Rx V6 Pkts        : N/A
Tx Pkts           : 15
Tx V4 Pkts        : 0
Tx V4 Discard Pkts: 0
Tx V6 Pkts        : 0
Tx V6 Discard Pkts: 0
Mpls Rx Pkts      : 0
Mpls Tx Pkts      : 0
Rx Bytes          : 62236
Rx V4 Bytes       : N/A
Rx V6 Bytes       : N/A
Tx Bytes          : 900
Tx V4 Bytes       : 0
Tx V4 Discard Byt*: 0
Tx V6 Bytes       : 0
Tx V6 Discard Byt*: 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           : 10.3.1.3*
Lease Populate    : 0
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

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
Time (seconds)   : - 10
Time (seconds)   : - 10
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.
*B:Dut-C#

arp

Syntax `arp [ip-address] | [mac ieee-address] | [sap sap-id] | [interface ip-int-name] [sdp sdp-id:vc-id]`

Context `show>service>id`

Description 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.

Parameters *ip-address* — Displays only ARP entries in the ARP table with the specified IP address.

Default All IP addresses.

mac ieee-address — Displays only ARP entries in the ARP table with the specified 48-bit MAC address. The MAC address can be expressed in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

Default All MAC addresses.

sap sap-id — Displays SAP information for the specified SAP ID. See Common CLI Command Descriptions on page 849 for command syntax.

port-id — **interface** — Specifies matching service ARP entries associated with the IP interface.

ip-address — The IP address of the interface for which to display matching ARP entries.

Values 1.0.0.0 — 223.255.255.255

ip-int-name — The IP interface name for which to display matching ARPs.

sdp-id — The SDP identifier.

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

Output **Show Service-ID ARP** — The following table describes show service-id ARP output fields.

Label	Description
IP Address	The IP address.
MAC Address	The specified MAC address.
Type	Static — FDB entries created by management. Learned — Dynamic entries created by the learning process. OAM — Entries created by the OAM process. Other — Local entries for the IP interfaces created.
Expiry	The age of the ARP entry.
Interface	The interface applied to the service.
SAP	The SAP ID.

Sample Output

base

Syntax base

Context show>service>id

Description This command displays basic information about this IES service.

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#
```

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.
ip-address — Specifies the ip-address of the neighbor.

Sample 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.

```
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.
- ip-address** — Specifies the ip-address of the neighbor.
- port** — Specifies the neighbor TCP port number use for this ANCP session.

Values 0 — 65535

association — Displays to what object the ANCP-string is associated.

statistics — Displays statistics information about an ANCP session known to the system.

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
=====
```

IES Show Commands

```
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#
```

host

Syntax host

Context show>service>id

Description Displays static hosts configured for this IES service.

Output **Show All Service-ID Output** — The following table describes the show all service-id command output fields.

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.

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   143.144.145.1  N/A              N/A
N/A                               Fwding
-----
Number of static hosts : 1
=====
*A:ALA-48#
```

host-connectivity-verify

Syntax `host-connectivity-verify statistics [sap sap-id]`

Context `show>service>id`

Description Displays host connectivity check statistics.

Parameters `statistics` — Displays host connectivity verification data.
`sap sap-id` — See Common CLI Command Descriptions on page 849 for command syntax.

Output **Show Service Id Host Connectivity Verify** — The following table describes show service-id host connectivity verification output fields:

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.

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  143.144.145.1          Up
=====
A:ALA-48>show>service>id#
```

interface

Syntax `interface [ip-address | ip-int-name] [interface-type] [detail] [family]`

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

- 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
 x: [0 — FFFF]H
 d: [0 — 255]

ip-int-name — Specifies the IP interface name for which to display information.

Values 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 to display either group or interfaces.

Values group, subscriber

detail — Displays detailed IP interface information.

Default IP interface summary output.

Output Show Service-ID — The following table describes show service-id output fields.

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.

Label	Description (Continued)
Redirects	Specifies the rate for ICMP redirect messages.
Unreachables	Specifies the rate for ICMP unreachable messages.
TTL Expired	Specifies the rate for ICMP TTL messages.

Sample Output

```
A:ALA-49# show service id 88 interface
=====
Interface Table
=====
Interface-Name          Adm      Opr (v4/v6)  Type      Port/SapId
  IP-Address                               PfxState
-----
Sector A                Up       Down/Down   IES       1/1/1.2.2
-
test                    Up       Down/Down   IES       1/1/2:0
  1.1.1.1/31            n/a
  1.1.1.1/31            n/a
  1.1.2.1/31            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#
```

labels

Syntax labels

Context show>service>id

Description Displays the labels being used by the service.

Output **Show Service-ID Labels** — The following table describes show service-id labels output fields:

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.

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#
```

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>sdp

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* — Specifies the SDP ID for which to display information.

Default All SDPs.

Values 1 — 17407

pw-port-id — Specifies the pseudo-wire port identifier.

Values 1 — 10239

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

Default SDPs with any far-end IP address.

detail — Displays detailed SDP information.

Default SDP summary output.

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

Default SDP summary output.

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. Octets.   : 0                I. Dro. Octets.   : 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

IES Show Commands

```

-----
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:Dut-B>config>service>sdp# show service sdp 230 detail

-----
Service Destination Point (Sdp Id : 230) 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-OSPF

Admin State                 : Up                          Oper State              : Up
Signaling                  : TLDP                       Metric                  : 0
Acct. Pol                  : None                       Collect Stats           : Disabled
Last Status Change        : 05/27/2015 06:59:46  Adv. MTU Over.         : No
Last Mgmt Change          : 05/27/2015 06:59:46  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-OSPF

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
-----
OSPF                : enabled          LSP Id            : 524291
Oper Instance Id    : 0

=====
*A:Dut-B>config>service>sdp#

```

pw-port

Syntax **pw-port** [*pw-port-id*] [**detail**]
pw-port sdp *sdp-id*
pw-port sdp none

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 — 10239

detail — Displays detailed port information that includes all the **pw-port** output fields.

sdp *sdp-id* — The SDP ID for which to display matching PW port information.

Values 1 — 17407

Output Show PW-Port — The following table describes **show pw-port** output fields:

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.

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

```
=====
```

sap

Syntax `sap sap-id [detail]`

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 *sap-id* — The ID that displays SAPs for the service in the *slot/mda/port[.channel]* format. See Common CLI Command Descriptions on page 849 for command syntax.

detail — Displays detailed information for the SAP.

Output **Show Service-ID SAP** — The following table describes show service SAP fields:

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. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapPipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, 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.
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.

Label	Description (Continued)
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, etc.
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.
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 (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, etc.
Ingress TD Profile	The profile ID applied to the ingress SAP.
Egress TD Profile	The profile ID applied to the egress SAP.
Mult Svc Site	Specifies the customer's multi-service-site name.
I. Sched Pol	The ingress scheduler policy applied to the customer's multi-service-site.
E. Sched Pol	The egress scheduler policy applied to the customer's multi-service-site.

Sample Output

```
A:ALA-49# show service id 88 sap 771/1/1.2.2
=====
Service Access Points(SAP)
```

IES Show Commands

```

=====
Service Id      : 88
SAP             : 1/1/1.2.2          Encap           : bcpNull
Admin State    : Up                 Oper State      : Down
Flags          : PortOperDown
                SapEgressQoSMismatch
Last Status Change : 06/06/2006 08:22:07
Last Mgmt Change  : 06/06/2006 14:15:58
Admin MTU       : 1518              Oper MTU        : 1518
Ingress qos-policy : 2              Egress qos-policy : 1020
Shared Q plcy   : default           Multipoint shared : Enabled
Ingress Filter-Id : n/a            Egress Filter-Id : n/a
tod-suite      : None
Multi Svc Site  : None
Acct. Pol      : None               Collect Stats    : Disabled
Anti Spoofing   : None              Nbr Static Hosts : 0
-----
Subscriber Management
-----
Admin State      : Down              MAC DA Hashing   : False
Def Sub-Profile  : None
Def SLA-Profile  : None
Sub-Ident-Policy : None
Subscriber Limit : 1
Single-Sub-Parameters
  Prof Traffic Only : False
  Non-Sub-Traffic   : N/A
=====
A:ALA-49#

*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#

*A:PE# show service id 1 sap 1/1/1:1 detail
=====

```

Service Access Points(SAP)

```

=====
Service Id      : 1
SAP            : 1/1/1:1
Description    : (Not Specified)
Admin State    : Up
Flags          : None
Multi Svc Site : None
Last Status Change : 01/29/2015 10:51:49
Last Mgmt Change  : 01/28/2015 11:48:21
Sub Type       : regular
Dot1Q Ethertype : 0x8100
Split Horizon Group: (Not Specified)

Et tree Root Leaf Tag: Disabled
Et tree Leaf AC      : Disabled
Max Nbr of MAC Addr: No Limit
Learned MAC Addr    : 0
OAM MAC Addr        : 0
Host MAC Addr       : 0
SPB MAC Addr        : 0
BGP EVPN Addr       : 0
Admin MTU           : 1518
Ingr IP Fltr-Id     : n/a
Ingr Mac Fltr-Id    : n/a
Ingr IPv6 Fltr-Id   : n/a
tod-suite           : None

Q Frame-Based Acct  : Disabled
ARP Reply Agent     : Disabled
Mac Learning        : Enabled
Mac Aging           : Enabled
BPDU Translation    : Disabled
L2PT Termination    : Disabled
Vlan-translation    : None

Acct. Pol           : None

Anti Spoofing       : None
Avl Static Hosts    : 0
Calling-Station-Id : n/a

Application Profile: None
Transit Policy      : None

Oper Group          : (none)
Host Lockout Plcy   : n/a
Lag Link Map Prof   : (none)
Cflowd             : Disabled
MCAC Policy Name    :
MCAC Max Unconst BW: no limit
MCAC In use Mand BW: 0
MCAC In use Opnl BW: 0
Use LAG port weight: no
Restr MacProt Src   : Disabled
Auto Learn Mac Prot: Disabled
Time to RetryReset  : never
Mac Move            : Blockable
Egr MCast Grp      :
Auth Policy         : None

Encap              : q-tag
Oper State         : Up
QinQ Ethertype     : 0x8100

Et tree Leaf Tag    : 0
Total MAC Addr     : 0
Static MAC Addr    : 0
DHCP MAC Addr      : 0
Intf MAC Addr      : 0
Cond MAC Addr      : 0
EVPN Static Addr   : 0
Oper MTU           : 1518
Egr IP Fltr-Id     : n/a
Egr Mac Fltr-Id    : n/a
Egr IPv6 Fltr-Id   : n/a
qinq-pbit-marking  : both
Egr Agg Rate Limit: max
Limit Unused BW    : Disabled
Host Conn Verify   : Disabled
Discard Unkwn Srce: Disabled
Mac Pinning        : Disabled

Collect Stats      : Disabled

Dynamic Hosts      : Enabled
Tot Static Hosts   : 0

Monitor Oper Grp   : (none)

MCAC Const Adm St  : Enable
MCAC Max Mand BW   : no limit
MCAC Avail Mand BW: unlimited
MCAC Avail Opnl BW: unlimited

Restr MacUnpr Dst  : Disabled
RestProtSrcMacAct  : Disable
Retries Left       : 3
Blockable Level    : Tertiary

```


Def Inter-Dest-Id : None
Def App-Profile : None
Sub-Ident-Policy : None

Subscriber Limit : 1
Single-Sub-Parameters
Prof Traffic Only : False
Non-Sub-Traffic : N/A

Sap Statistics

Last Cleared Time : N/A

	Packets	Octets
CPM Ingress	: 0	0

Forwarding Engine Stats

Dropped	: 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 30)

Dro. InProf	: 0	0
Dro. OutProf	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

Sap per Queue stats

	Packets	Octets
--	---------	--------

Ingress Queue 1 (Unicast) (Priority)

Off. HiPrio	: 0	0
Off. LowPrio	: 0	0
Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

Ingress Queue 11 (Multipoint) (Priority)

Off. HiPrio	: 0	0
Off. LowPrio	: 0	0
Off. Managed	: 0	0
Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

Egress Queue 1

For. InProf	: 0	0
For. OutProf	: 0	0
Dro. InProf	: 0	0
Dro. OutProf	: 0	0

```
=====
*A:PE#
```

sdp

Syntax **sdp** [{*sdp-id* | **far-end** *ip-address* }] [**detail**]

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* — The SDP ID for which to display information.

Values 1 — 17407

far-end *ip-address* — When specified, displays SDP having the specified far-end IP address.

detail — Displays detailed information for the SDP.

Sample Output

```
A:Dut-A# show service id 1 sdp detail
=====
Services: Service Destination Points Details
=====
  Sdp Id 1:1  - (10.20.1.2)
-----
Description      : Default sdp description
SDP Id           : 1:1                               Type           : Spoke
VC Type          : Ether                             VC Tag         : n/a
Admin Path MTU   : 0                                 Oper Path MTU  : 9186
Far End          : 10.20.1.2                         Delivery       : MPLS

Admin State      : Up                               Oper State     : Up
Acct. Pol        : None                             Collect Stats  : Disabled
Ingress Label    : 2048                             Egress Label   : 2048
Ing mac Fltr     : n/a                               Egr mac Fltr  : n/a
Ing ip Fltr      : n/a                               Egr ip Fltr   : n/a
Ing ipv6 Fltr    : n/a                               Egr ipv6 Fltr : n/a
Admin ControlWord : Not Preferred                    Oper ControlWord : False
Last Status Change : 05/31/2007 00:45:43             Signaling      : None
Last Mgmt Change  : 05/31/2007 00:45:43

Class Fwding State : Up
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit                       Total MAC Addr : 0
Learned MAC Addr   : 0                               Static MAC Addr : 0

MAC Learning       : Enabled                         Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
L2PT Termination   : Disabled                       BPDU Translation : Disabled
MAC Pinning        : Disabled
```

```

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

Statistics            :
I. Fwd. Pkts.         : 0                    I. Dro. Pkts.       : 0
I. Fwd. Octs.         : 0                    I. Dro. Octs.       : 0
E. Fwd. Pkts.         : 0                    E. Fwd. Octets      : 0
MCAC Policy Name      :
MCAC Max Unconst BW   : no limit            MCAC Max Mand BW    : no limit
MCAC In use Mand BW   : 0                   MCAC Avail Mand BW  : unlimited
MCAC In use Opnl BW   : 0                   MCAC Avail Opnl BW  : unlimited

Associated LSP LIST :
Lsp Name              : A_B_1              Oper State           : Up
Admin State           : Up                 Time Since Last Tr* : 00h26m35s

Lsp Name              : A_B_2              Oper State           : Up
Admin State           : Up                 Time Since Last Tr* : 00h26m35s

Lsp Name              : A_B_3              Oper State           : Up
Admin State           : Up                 Time Since Last Tr* : 00h26m34s

Lsp Name              : A_B_4              Oper State           : Up
Admin State           : Up                 Time Since Last Tr* : 00h26m34s

Lsp Name              : A_B_5              Oper State           : Up
Admin State           : Up                 Time Since Last Tr* : 00h26m34s

Lsp Name              : A_B_6              Oper State           : Up
Admin State           : Up                 Time Since Last Tr* : 00h26m34s

Lsp Name              : A_B_7              Oper State           : Up
Admin State           : Up                 Time Since Last Tr* : 00h26m34s

Lsp Name              : A_B_8              Oper State           : Up
Admin State           : Up                 Time Since Last Tr* : 00h26m35s

Lsp Name              : A_B_9              Oper State           : Up
Admin State           : Up                 Time Since Last Tr* : 00h26m34s

Lsp Name              : A_B_10             Oper State           : Up
Admin State           : Up                 Time Since Last Tr* : 00h26m34s
-----
Class-based forwarding :
-----
Class forwarding      : enabled
Default LSP           : A_B_10             Multicast LSP        : A_B_9
=====

```

IES Show Commands

```
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#
```

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*

Context show>router>vrrp

Description This command displays statistics for the VRRP instance.

Parameters *interface-name* — Displays statistics for the specified interface.
virtual-router-id — Displays statistics for the specified virtual router ID.
Values 1 — 255

statistics

Syntax **statistics**
Context show>router>vrrp
Description This command displays statistics for the VRRP instance.

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.

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.

arp-host

Syntax **arp-host**
arp-host { **mac** *ieee-address* | **sap** *sap-id* | **ip-address** *ip-address[/mask]* }
arp-host [**port** *port-id*] [**inter-dest-id** *intermediate-destination-id* | **no-inter-dest-id**]
arp-host statistics [**sap** *sap-id* | **interface** *interface-name*]
Context clear>service>id
Description This command clears ARP host data.

interface

IES Show Commands

Syntax `interface [ip-int-name | ip-addr] [icmp]`

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 `ip-int-name | ip-addr` — The IP interface name or IP interface address.

Default All IP interfaces.

icmp — Specifies to reset the ICMP statistics for the IP interface(s) used for ICMP rate limit.

interface

Syntax `interface interface-name [vrid virtual-router-id]`
`interface interface-name vrid virtual-router-id ipv6`

Context `clear>router>vrrp`

Description This command clears and resets VRRP instances.

Parameters `interface-name` — Specifies an existing interface name up to 32 characters in length.

`virtual-router-id` — Specifies the virtual router identifier.

Values 1 — 255

statistics

Syntax `statistics interface interface-name [vrid virtual-router-id]`
`statistics`
`statistics interface interface-name vrid virtual-router-id ipv6`

Context `clear>router>vrrp`

Description This command clears statistics for VRRP instances.

Parameters `interface-name` — Specifies an existing interface name up to 32 characters in length.

`virtual-router-id` — Specifies the virtual router identifier.

Values 1 — 255

fdb

Syntax **fdb** {**all** | **mac** *ieee-address* | **sap** *sap-id*] | **mesh-sdp** *sdp-id[:vc-id]* | **spoke-sdp** *sdp-id:vc-id*}

Context clear>service>id

Description This command clears FDB entries for the service.

Parameters **all** — Clears all FDB entries.

mac *ieee-address* — Clears only FDB entries in the FDB table with the specified 48-bit MAC address. The MAC address can be expressed in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

sap *sap-id* — Clears the specified SAP information. See Common CLI Command Descriptions on page 849 for command syntax.

mesh-sdp — Clears only service FDB entries associated with the specified mesh SDP ID. For a mesh SDP, the VC ID is optional.

spoke-sdp — Clears only service FDB entries associated with the specified spoke SDP ID. For a spoke SDP, the VC ID must be specified.

sdp-id — The SDP ID for which to clear associated FDB entries.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to clear associated FDB entries.

Default For mesh SDPs only, all VC IDs.

1 — 4294967295site

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 in length.

Values 1 — 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 — 17407

IES Show Commands

vc-id — The virtual circuit ID on the SDP ID to be reset.

Values 1 — 4294967295

stp

Syntax **stp**

Context clear>service>statistics>id

Description Clears all spanning tree statistics for the service ID.

IES Debug Commands

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** *ip-address*

Context debug>service>id>host-connectivity-verify

Description This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular IP address.

Parameters *ip-address* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

mac

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

Context debug>service>id>host-connectivity-verify

Description This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular MAC address.

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.

sap

Syntax [no] sap *sap-id*

Context debug>service>id>host-connectivity-verify

Description This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular SAP.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 849 for command syntax.

packets

Syntax [no] packets
 [no] packets interface *ip-int-name* [vrid *virtual-router-id*]
 [no] packets interface *ip-int-name* vrid *virtual-router-id* ipv6

Context debug>router>vrrp

Description This command enables or disables debugging for VRRP packets.

Parameters *ip-int-name* — Specifies the interface name.
virtual-router-id — Specifies the router ID.

Values 1 — 255

events

Syntax [no] events
 [no] events interface *ip-int-name* [vrid *virtual-router-id*]
 [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 — 255

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.

Virtual Private Routed Network Service

In This Chapter

This chapter provides information about the Virtual Private Routed Network (VPN) service and implementation notes.

Topics in this chapter include:

- [VPRN Service Overview on page 225](#)
- [VPRN Features on page 241](#)
 - [IP Interfaces on page 241](#)
 - [Object Grouping and State Monitoring on page 243](#)
 - [SAPs on page 245](#)
 - [QoS Policies on page 246](#)
 - [Filter Policies on page 246](#)
 - [DSCP Marking on page 247](#)
 - [Configuration of TTL Propagation for VPRN Routes on page 249](#)
 - [CE to PE Routing Protocols on page 251](#)
 - [PE to PE Tunneling Mechanisms on page 251](#)
 - [Per VRF Route Limiting on page 251](#)
 - [Spoke SDPs on page 252](#)
 - [IP-VPNs on page 254](#)
 - [IPCP Subnet Negotiation on page 255](#)
 - [Cflowd for IP-VPNs on page 256](#)
 - [Carrier Supporting Carrier \(CsC\) on page 257](#)

(Continued)

- [Traffic Leaking to GRT on page 260](#)
- [Traffic Leaking from VPRN to GRT for IPv6 on page 261](#)
- [RIP Metric Propagation in VPRNs on page 262](#)
- [NTP Within a VPRN Service on page 263](#)
- [PTP Within a VPRN Service on page 263](#)
- [VPN Route Label Allocation on page 264](#)
- [Multicast in IP-VPN Applications on page 268](#)
 - [Use of Data MDTs on page 270](#)
 - [Multicast Protocols Supported in the Provider Network on page 271](#)
 - [MVPN Membership Auto-discovery using BGP on page 271](#)
 - [MVPN \(Rosen\) Membership Auto-Discovery using BGP MDT-SAFI on page 274](#)
 - [PE-PE Transmission of C-Multicast Routing using BGP on page 274](#)
 - [VRF Route Import Extended Community on page 274](#)
 - [Provider Tunnel Support on page 276](#)
 - [MVPN \(NG-MVPN\) Upstream Multicast Hop Fast Failover on page 285](#)
 - [Multicast VPN Extranet on page 286](#)
 - [IPv6 MVPN Support on page 293](#)
 - [Multicast Core Diversity for Rosen MDT_SAFI MVPNs on page 296](#)
 - [NG-MVPN Multicast Source Geo-Redundancy on page 298](#)
 - [Multicast Core Diversity for Rosen MDT SAFI MVPNs on page 301](#)
 - [Inter-AS MVPN on page 303](#)
- [Configuring a VPRN Service with CLI on page 309](#)
- [Common Configuration Tasks on page 312](#)
- [Service Management Tasks on page 326](#)

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 7](#) displays a VPRN network diagram example.



OSSG024

Figure 7: Virtual Private Routed Network

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)

Core MP-BGP Support

BGP is used with BGP extensions mentioned in [Routing Prerequisites on page 227](#) 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.

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 7950 XRS implementation supports the three (3) **Type** values as defined in the standard.




Figure 8: 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.

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 9](#) 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 9: Basic eBGP Topology

Another example displayed in [Figure 10](#) 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 eBGP is enabled) as if only a single VRF was involved.




Figure 10: 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.

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.

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.

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.

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 VPRN 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.

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 SR-Series router will continue to send polls and re-activate any routes that are restored.

Constrained Route Distribution (RT Constraint)

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.

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.

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`
- `config>service>vpls>bgp`

Constrained Route Distribution (RT Constraint)

- `config>service>epipe>bgp`

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, , or 7950 XRS 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.

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.

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. Note that 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

Constrained Route Distribution (RT Constraint)

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.

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 9](#).

Note that BGP fast reroute information specific to the base router BGP context is described in the BGP Fast Reroute section of the 7x50 SR OS Routing Protocols Guide.

Table 9: 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, but if the VPN-IP routes are label-per-prefix the ingress card must be FP2 or better
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, but if the VPN-IP routes are label-per-prefix the ingress card must be FP2 or better for PIC
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, but if the VPN-IP routes are label-per-prefix the ingress card must be FP2 or better
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 9: BGP Fast Reroute Scenarios (VPRN Context)

Ingress Packet	Primary Route	Backup Route	Prefix Independent Convergence
MPLS (egress)	IPv6 route with next-hop A resolved by an IPv6 route	Yes, but if the VPN-IP routes are label-per-prefix the ingress card must be FP2 or better for PIC	VPRN label mode must be VRF. VPRN must export its VPN-IP routes with RD ≠ y. For the best performance the backup next-hop must advertise the same VPRN label value with all routes (e.g. per VRF label).

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 BGP Fast Reroute section of the 7750 SR OS Routing Protocols Guide for more information.

VPRN Features

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

- [IP Interfaces on page 241](#)
 - [SAPs on page 245](#)
 - [Encapsulations on page 245](#)
 - [QoS Policies on page 246](#)
 - [Filter Policies on page 246](#)
 - [CE to PE Routing Protocols on page 251](#)
 - [PE to PE Tunneling Mechanisms on page 251](#)
 - [Per VRF Route Limiting on page 251](#)
 - [Using OSPF in IP-VPNs on page 254](#)
 - [Spoke SDPs on page 252](#)
 - [Carrier Supporting Carrier \(CsC\) on page 257](#)
 - [Multicast in IP-VPN Applications on page 268](#)
-

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

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 \(QPPB\) on page 18](#) section on page 15 and the IP Router Configuration section in the 7x50 OS Router Configuration Guide.

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.

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:

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 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 example uses an R-VPLS context. The VPLS instance includes the **allow-ip-int-bind** and the **service-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 customer 1 create
  allow-ip-int-bind
  stp
    shutdown
  exit
  service-name "v1"
  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

vprn 2001 customer 1 create
  interface "i2001" create
    address 21.1.1.1/24
    monitor-oper-group "g1"
    vpls "v1"
  exit
  no shutdown
  exit
```

SAPs

Encapsulations

The following SAP encapsulations are supported on the SR-OS VPRN service:

- Ethernet null
- Ethernet dot1q

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.

Note that both Layer 2 or Layer 3 criteria can be used in the QoS policies for traffic classification in an VPRN.

Filter Policies

Ingress and egress IPv4 and IPv6 filter policies can be applied to VPRN SAPs.

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 linecard. The egress linecard 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 10: 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
IGMP/MLD	Yes	Yes	Yes	Yes	AF
Telnet	Yes	Yes	Yes	Yes	AF
TFTP	Yes		Yes	Yes	AF
FTP	Yes				AF

Table 10: DSCP/FC Marking (Continued)

Protocol	IPv4	IPv6	DSCP Marking	Dot1P Marking	Default FC
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
Bootp	Yes				AF
IPv6 Neighbor Discovery	Yes				NC

Default DSCP Mapping Table

DSCP Name	DSCP Value	DSCP Value	DSCP Value	Label
	Decimal	Hexadecimal	Binary	

Default	0	0x00	0b000000	be
nc1	48	0x30	0b110000	h1
nc2	56	0x38	0b111000	nc
ef	46	0x2e	0b101110	ef
af11	10	0x0a	0b001010	assured
af12	12	0x0c	0b001100	assured
af13	14	0x0e	0b001110	assured
af21	18	0x12	0b010010	11
af22	20	0x14	0b010100	11
af23	22	0x16	0b010110	11
af31	26	0x1a	0b011010	11
af32	28	0x1c	0b011100	11
af33	30	0x1d	0b011110	11
af41	34	0x22	0b100010	h2
af42	36	0x24	0b100100	h2
af43	38	0x26	0b100110	h2
default*0				

*The default forwarding class mapping is used for all DSCP names/values for which there is no explicit forwarding class mapping.

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.

Configuration of TTL Propagation for VPRN Routes

- 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]`

Note 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.

Note that 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.

CE to PE Routing Protocols

The SR-OS VPRN supports the following PE to CE routing protocols:

- BGP
 - Static
 - RIP
 - OSPF
-

PE to PE Tunneling Mechanisms

The SR-OS supports multiple mechanisms to provide transport tunnels for the forwarding of traffic between PE routers within the 2547bis network.

The SR-OS VPRN implementation supports the use of:

- RSVP-TE protocol to create tunnel LSP's between PE routers
- LDP protocol to create tunnel LSP's between PE routers
- 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 LSP's with bandwidth reservations for specific customers as is available with explicit SDPs for the service.

Per VRF Route Limiting


The SR-OS allows setting the maximum number of routes that can be accepted in the VRF for a VPRN service. There are options to specify a percentage threshold at which to generate an event that the VRF table is near full and an option to disable additional route learning when full or only generate an event.

Spoke SDPs

Distributed services use service distribution points (SDPs) to direct traffic to another SR-Series router via service tunnels. SDPs are created on each participating SR-Series and then bound to a specific service. SDP can be created as either GRE or MPLS. Refer to the *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 11 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.



al_0163

Figure 11: SDP-ID and VC Label Service Identifiers

T-LDP Status Signaling for Spoke-SDPs Terminating on IES/VRPN

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 VRPN 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 VRPN 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 VRPN 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/VRPN PE.

This feature can be used to provide redundant connectivity to a VRPN or IES from a PE providing a VLL service, as shown in [Figure 12](#).





Figure 12: Active/Standby VRF Using Resilient Layer 2 Circuits



IP-VPNs

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.

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 13](#).




Figure 13: 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.

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.

Carrier Supporting Carrier (CsC)

Carrier Supporting Carrier (CSC) is a solution 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 SROS allows a 7x50 to be deployed as any of the following devices shown in Figure 14:

- 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 14: Carrier Supporting Carrier Reference Diagram

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.

Carrier Supporting Carrier (CsC)

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.

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.

Traffic Leaking to GRT

Traffic leaking to Global Route Table (GRT) allows service providers to offer VPRN and Internet services to their customers over a single VRF interface. This currently supports IPv4 .

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”, under the “enable-grt” context. 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

Traffic Leaking from VPRN to GRT for IPv6

This feature allows IPv6 destination lookups in two distinct routing tables. 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.

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 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 15: RIP Metric Propagation in VPRNs

NTP Within a VPRN Service

The NTP within a VPRN service enables the service router to act as the NTP server to CPE devices on a VPRN interface. Individual VPRN interfaces may be configured to listen to and respond to client requests, or additionally may be configured to send NTP broadcast messages. Authentication keys 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.

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 Basic System Configuration Guide.

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.

[Table 11](#) summarizes the logic that determines the label allocation method for an exported route.

Table 11: Label allocation logic

1	If the IP route is LOCAL, AGGREGATE, or BGP-VPN always use the VRF label.
2	If the IP route is accepted by a VRF export policy with the advertise-label per-prefix action, use LPP.
3	If the IP (BGP) route is exported by the export-inactive-bgp command (VPRN best external), use LPN.
4	If the IP route is exported by a VPRN configured for label-mode next-hop, use LPN.
5	Else, use the per-VRF label.

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 16](#).




Figure 16: 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 multicast-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. The above command is ignored on FP1-based hardware.

Multicast in IP-VPN Applications

This section and its subsections focuses on Multicast in IP VPN functionality. A reader should familiarize itself first with Multicast section in SROS Routing Protocols Guide where multicast protocols (PIM, IGMP, MLD, MSDP) are described.

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 17: Multicast in IP-VPN Applications

Figure 17 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.

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.

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.

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 12 and Table 13 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 12: Supported Configuration Combinations

Auto-Discovery	Inclusive PIM SSM	Action
Yes	Yes	Allowed
No	Yes	Not Allowed
Yes or No	No	Allowed

Table 13: Supported Configuration Combinations

Auto-Discovery	C-Mcast-Signaling	s-PMSI A-D	Action
Yes	BGP	Ignored	Allowed
Yes	PIM	Yes	Allowed
Yes	PIM	No	Allowed
No	BGP	Ignored	Not Allowed
No	PIM	Ignored	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**, S-PMSI A-D configuration value is used.

MVPN (Rosen) Membership Auto-Discovery using BGP MDT-SAFI

MVPN implementation based on the draft *-Rosen* 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.

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.

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 vrftraget 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;
```

Provider Tunnel Support

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. For more details, see [Wildcard \(C-*, C-*\) P2MP LSP S-PMSI](#). 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 details, 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.

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.

OPERATIONAL NOTE: All SR OS routers 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.

OPERATIONAL 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, SROS implements two modes of operations for BSR.

By default (bsr unicast):

Provider Tunnel Support

- 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 IPMSI 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.

OPERATIONAL NOTES:

- 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).

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.

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 18 illustrates dynamic mLDP signaling for IP multicast in VPRN.




Figure 18: Dynamic mLDP Signaling for IP Multicast in VPRN

As illustrated in Figure 18, 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.
- The feature requires chassis mode C.

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 to optimize 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 19](#) depicts 4-site MVPN with sender-only, receiver-only and sender-receiver (default) sites.




Figure 19: 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.

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-PSMI). 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)). Note that 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.

Note that 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.

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.
-

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.

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.

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:

- Multicast extranet functionality requires all network ports to be on FP2 or newer line cards.
- 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).

Multicast VPN Extranet

- 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). Note that 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:

- Multicast extranet functionality requires all network ports to be on FP2 or newer line cards.
- 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).

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 21](#) represents a sample deployment model.




Figure 21: Source PE Transit Replication and Receiver PE Per-group SSM Extranet Mapping

The architecture displayed in [Figure 21](#) 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 subsections 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.
- Per-group mapping requires network interfaces to be on FP2-based or newer line cards.

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 22](#) depicts a sample deployment.




Figure 22: 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.
- The feature requires FP2-based line cards or newer for all network ports.

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 23 depicts PIM ASM extranet map support.




Figure 23: 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 (`configure>service>vprn>mvpn>rpf-select>core-mvpn` or `configure>service>vprn>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.
- This feature requires network chassis mode D.

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 RSVPE-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 24](#) shows an example of an operator with IPv4 MPLS backbone providing IPv6 MVPN service to Customer 1 and Customer 2.



al_0168

Figure 24: IPv6 MVPN Example

SROS 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:

1. IPv6 MVPN requires chassis mode D
2. Non-congruent topologies are not supported
3. IPv6 is not supported in MCAC
4. 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.
5. IPv6 Auto-RP is not supported

Multicast Core Diversity for Rosen MDT_SAFI MVPNs

Figure 25 depicts Rosen MVPN core diversity deployment:



al_0485

Figure 25: 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. Note that 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-VPRN.

On source PEs (PE1: UMH, PE2: UMH in Figure 25), 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 SROS Release 12.0.R4 - BGP Connector support for non-default instances).

The above configuration ensures that MDT SAFI and IPVPN 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.

Note that 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. Note that 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.
- The feature requires all ports to be present on IOM3-XP or newer (Chassis Mode D). The restriction is not enforced. Failing to observe this restriction will prevent the feature to operate properly in the network.

NG-MVPN Multicast Source Geo-Redundancy

Multicast source geo-redundancy is targeted primarily for MVPN deployments for multicast delivery services like IPTV. The solutions 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-IPMSI 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 26 shows an operational example of multicast source geo-redundancy:




Figure 26: 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).
- Note: 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**, and **unicast-rt-pref** (no support for tunnel-status).

Multicast Core Diversity for Rosen MDT SAFI MVPNs

Figure 27 shows a Rosen MVPN core diversity deployment.



al_0379

Figure 27: Multicast Core Diversity

Core diversity allows operators 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 the main feature attributes:

- Rosen MVPN IPv4 multicast with MDT SAFI is supported with default and data MDTs.
- Rosen MVPN can use a non-default OSPF or ISIS instance (using their loopback addresses instead of a system address).
- Up to 3 distinct core instances are supported: system + 2 non-default OSPF instances shown in Figure 27.
- Note that 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.
- The feature is supported with CSC-VPRN.

Multicast Core Diversity for Rosen MDT SAFI MVPNs

On source PEs (PE1: UMH, PE2: UMH in the above picture), an MVPN is assigned to a non-default IGP core instance as follows:

- MVPN is statically pointed to use one of the non-default IGP instances loopback addresses as source address instead of system loopback IP.
- 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 IPVPN 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. Note that 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. Note that 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 R12.0R4) 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.
- The feature requires all ports to be present on IOM3-XP or newer (Chassis Mode D). The restriction is not enforced. Failing to observe this restriction will prevent the feature to operate properly in the network.

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.

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.

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]

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 28](#) depict set-up of a default MDT:




Figure 28: Inter-AS Option B Default MDT Setup

SROS 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 SROS implementation was designed also to interoperate with older routers Inter-AS implementations that do not comply with the RFC 5384 and RFC 5496.

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 29 depicts a default MDT setup:




Figure 29: 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 SROS 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 SROS router (if supported) using **configure router pim rpfv core mvpn**: PIM joins received can be a mix of core and mvpn RPF vectors.

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.

Configuring a VPRN Service with CLI

This section provides information to configure Virtual Private Routed Network (VPRN) services using the command line interface.

Topics in this section include:

- [Basic Configuration on page 310](#)
- [Common Configuration Tasks on page 312](#)
 - [Configuring VPRN Components on page 313](#)
 - [Creating a VPRN Service on page 313](#)
 - [Configuring Global VPRN Parameters on page 314](#)
 - [Configuring VPRN Log Parameters on page 315](#)
 - [Configuring VPRN Protocols - BGP on page 318](#)
 - [Configuring VPRN Protocols - OSPF on page 322](#)
 - [Configuring a VPRN Interface on page 323](#)
 - [Configuring a VPRN Interface SAP on page 324](#)
- [Service Management Tasks on page 326](#)
 - [Modifying VPRN Service Parameters on page 326](#)
 - [Deleting a VPRN Service on page 327](#)
 - [Disabling a VPRN Service on page 328](#)
 - [Re-enabling a VPRN Service on page 330](#)

Basic Configuration

The following fields require specific input (there are no defaults) to configure a basic VPRN service:

- Customer ID (refer to the *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 11.1.0.1/24
  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 6.5.0.0/24 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#
```

Common Configuration Tasks

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

1. Associate a VPRN service with a customer ID.
2. Define an autonomous system (optional).
3. Define a route distinguisher (mandatory).
4. Define VRF route-target associations or VRF import/export policies.
5. Define PIM parameters (optional).
6. Create an interface.
7. 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).
8. Define BGP parameters (optional).
 - BGP must be enabled in the `config>router>bgp` context.
9. Define RIP parameters (optional).
10. Define spoke SDP parameters (optional).
11. Create confederation autonomous systems within an AS. (optional).
12. Enable the service.

Configuring VPRN Components

This section provides VPRN configuration examples for the following entities:

- [Creating a VPRN Service on page 313](#)
 - [Configuring Global VPRN Parameters on page 314](#)
 - [Configuring Router Interfaces on page 317](#)
 - [Configuring VPRN Protocols - OSPF on page 322](#)
 - [Configuring a VPRN Interface SAP on page 324](#)
 - [Configuring VPRN Protocols - BGP on page 318](#)
-

Creating a VPRN Service

Use the following CLI syntax to create a VPRN service. A route distinguisher must be defined in order for VPRN to be operationally active.

CLI Syntax: `config>service# vprn service-id [customer customer-id]
 route-distinguisher [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#
```

Configuring Global VPRN Parameters

Refer to [VPRN Services Command Reference on page 331](#) 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#
```

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 "gkYL94190FFgu91PiRNvn3Rn10edkMU1" 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 3
          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#
```

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
```

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#
```

Configuring Router Interfaces

Refer to the 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 2.2.2.1/24
        port 1/1/33
    exit
    interface "if2"
        address 10.49.1.46/24
        port 1/1/34
    exit
    interface "if3"
        address 11.11.11.1/24
        port 1/1/35
    exit
...
#-----
ALA48>config>router#
```

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 on page 314](#).
- Specify a router ID - Note that if a new or different router ID value is entered in the BGP context, then the new value takes precedence and overwrites the VPRN-level router ID. See [Configuring Global VPRN Parameters on page 314](#).
- 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:	<code>config>service>vprn>bgp#</code>	(global level)
	<code>group</code>	(group level)
	<code>neighbor</code>	(neighbor level)

Note that the local-address must be explicitly configured if two systems have multiple BGP peer sessions between them for the session to be established.

For more information about the BGP protocol, refer to the OS 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 ([BGP Configuration Commands on page 356](#)).

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 11.1.0.1/24
        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 6.5.0.0/24 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#
```

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 SR-OS Router Configuration Guide.

CLI Syntax: `config>service>vprn>ospf#`

VPRN OSPF CLI Syntax

Refer to [OSPF Configuration Commands on page 365](#) 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 SR-OS Router Configuration Guide.

Configuring a VPRN Interface

Interface names associate an IP address to the interface, and then associate the IP interface with a physical port. The logical interface can associate attributes like an IP address, port, Link Aggregation Group (LAG) or the system.

There are no default interfaces.

Note that you can configure a VPRN interface as a loopback interface by issuing the `loopback` command instead of the `sap sap-id` command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

When using `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 Configuration Commands on page 365](#) 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 11.1.0.1/24
        exit
      exit
      static-route 6.5.0.0/24 next-hop 10.1.1.2
      spoke-sdp 2 create
      exit
      no shutdown
    exit
...
-----
*A:ALA-1>config>service#
```

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.

Refer to [OSPF Configuration Commands on page 365](#) 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 11.1.0.1/24
        sap 1/1/10:1 create
          ingress
            scheduler-policy "SLA2"
            qos 100
          exit
          egress
            scheduler-policy "SLA1"
            qos 1010
            filter ip 6
          exit
        exit
      static-route 6.5.0.0/24 next-hop 10.1.1.2
      spoke-sdp 2 create
      exit
      no shutdown
    exit
  ...
-----
*A:ALA-1>config>service#
```

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
  ipsec-interface "ipsec-private" create
    sap ipsec-1.private:1 create
    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
    no shutdown
    exit
  exit
  exit
  interface "corporate-network" create
    address 172.17.118.118/24
    sap 1/1/2 create
  exit
  exit
  static-route 172.16.91.0/24 ipsec-tunnel "remote-office"
  no shutdown
  exit
  exit
...
-----
*A:ALA-49>config#
```

Service Management Tasks

This section discusses the following service management tasks:

- [Modifying VPRN Service Parameters on page 326](#)
- [Deleting a VPRN Service on page 327](#)

Modifying VPRN Service Parameters

Use the CLI syntax to modify VPRN parameters ([VPRN Services Command Reference on page 331](#)).

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 6.5.0.0/24 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#
```

Deleting a VPRN Service

An 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
```

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 11.1.0.1/24
    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 6.5.0.0/24 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
exit
...
```

*A:ALA-1>config>service#

Re-enabling a VPRN Service

To re-enable a VPRN service that was shut down.

CLI Syntax: `config>service#
vprn service-id [customer customer-id]
no shutdown`

VPRN Services Command Reference

Command Hierarchies

- [VPRN Service Configuration Commands on page 332](#)
 - [IGMP Commands on page 335](#)
 - [Log Commands on page 337](#)
 - [Multicast VPN Commands on page 339](#)
 - [Interface Commands on page 343](#)
 - [Interface Spoke SDP Commands on page 348](#)
 - [Interface SAP Commands on page 352](#)
 - [Interface VRRP Commands on page 350](#)
 - [Oper Group Commands on page 355](#)
 - [Network Interface Commands on page 346](#)
 - [Network Ingress Commands on page 355](#)
 - [BGP Configuration Commands on page 356](#)
 - [ISIS Configuration Commands on page 361](#)
 - [OSPF Configuration Commands on page 365](#)
 - [PIM Configuration Commands on page 368](#)
- [Show Commands on page 372](#)
- [Clear Commands on page 374](#)
- [Debug Commands on page 375](#)

VPRN Service Configuration Commands

```

config
  — service
    — vprn service-id [customer customer-id]
    — no vprn service-id
      — aggregate ip-prefix/ip-prefix-length [summary-only] [as-set] [aggregator as-number:ip-address] [black-hole] [community comm-id] [description description]
      — aggregate ip-prefix/ip-prefix-length [summary-only] [as-set] [aggregator as-number:ip-address] [community comm-id] [indirect ip-address] [description description]
      — no aggregate ip-prefix/ip-prefix-length
      — allow-export-bgp-vprn
      — no allow-export-bgp-vprn
      — auto-bind-tunnel
        — resolution {any | filter | disabled }
        — [no] resolution-filter gre
        — [no] resolution-filter ldp
        — [no] resolution-filter rsvp
        — [no] resolution-filter sr-isis
        — [no] resolution-filter sr-ospf
      — autonomous-system as-number
      — no autonomous-system
      — backup-path [ipv4][ipv6]
      — [no] carrier-carrier-vpn
      — confederation confed-as-num members as-number [as-number...(up to 15 max)]
      — no confederation confed-as-num members as-number [as-number...(up to 15 max)]
      — 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
      — eth-cfm
        — tunnel-fault [accept | ignore]
      — grt-lookup
        — [no] enable-grt
          — [no] allow-local-management
          — static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [enable|disable] grt
          — no static-route
        — export-grt policy-name [policy-name ...(up to 5 max)]
        — export-limit num-routes

```

- **no export-limit**
- **export-v6-limit**
- **no export-v6-limit**
- **[no] hash-label**
- **igmp-host-tracking**
 - **expiry-time** *expiry-time*
 - **no expiry-time**
 - **[no] shutdown**
- **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**
- **network**
 - **ingress**
 - **filter** {**ip** *ip-filter-id* | **ipv6** *ipv6-filter-id*}
 - **no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
- **[no] ptp**
 - **peer** *a.b.c.d* [**create**]
 - **no peer** *a.b.c.d*
 - **[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**
- **service-name** *service-name*
- **no service-name**
- **sgt-qos**
 - **application** *dscp-app-name* **dscp** {*dscp-value* | *dscp-name*}
 - **application** *dot1p-app-name* **dot1p** *dot1p-priority*
 - **no application** {*dscp-app-name* | *dot1p-app-name*}
 - **dscp** *dscp-name* **fc** *fc-name*
 - **no dscp** *dscp-name*
- **[no] shutdown**
- **snmp**
 - **[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*
- **[no] spoke-sdp** *sdp-id*
 - **[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** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**community** *comm-id*] [**enable** | **disable**] {**next-hop** *ip-int-name*|*ip-address* [**mcast-family**] | **ipsec-tunnel** *ipsec-tunnel-name*} [**bfd-enable** | {**cpe-check** *cpe-ip-address* [**interval** *seconds*] [**drop-count** *count*] [**padding-size** *padding-size*] [**log**]}
- **[no] static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**community** *comm-id*] [**enable** | **disable**] **indirect** *ip-address* [**cpe-check** *cpe-ip-address* [**interval** *seconds*]] [**drop-count** *count*] [**padding-size** *padding-size*]] [**log**]
- **[no] static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**community** *comm-id*] [**enable** | **disable**] **black-hole** [**mcast-family**]
- **ttl-propagate**
 - **local** [**inherit**|**none**|**vc-only**|**all**]
 - **transit** [**inherit**|**none**|**vc-only**|**all**]
- **type** {**hub** | **spoke** | **subscriber-split-horizon**}
- **no type**
- **vrf-export** *policy-name* [*policy-name...*(upto 16 max)]
- **no vrf-export**
- **vrf-import** *policy-name* [*policy-name...*(upto 16 max)]
- **no vrf-import**
- **vrf-target** {*ext-comm* | { [**export** *ext-comm*] [**import** *ext-comm*] }}
- **no vrf-target**

IGMP Commands

```

config
  — service
    — vprn service-id [customer customer-id]
    — no vprn 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 [1..32000]
          — no max-grp-sources
          — max-sources [1..1000]
          — no max-sources
          — mcac
            — 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 [1..32000]
          — no max-sources
          — max-grp-sources [1..1000]
          — no max-grp-sources
          — mcac
            — 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** *grp-ip-address*
 - **[no] source** *ip-address*
 - **[no] starg**
- **[no] subnet-check**
- **version** *version*
- **no version**
- **[no] query-interval**
- **query-interval** *seconds*
- **[no] query-last-member-interval**
- **query-last-member-interval** *seconds*
- **[no] query-response-interval**
- **query-response-interval** *seconds*
- **[no] robust-count**
- **robust-count** *robust-count*
- **[no] shutdown**
- **ssm-translate**
 - **[no] grp-range** *start end*
 - **[no] source** *ip-address*

Log Commands

Refer to the 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**

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 [persistend-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> ...up to 8 maximum]
          — no src-prefix ip-address/mask
          — ipv6
            — src-prefix ipv6-ip-address/prefix-length [ipv6-ip-address/prefix-length > ...up to 8 maximum]
            — no src-prefix ipv6-ip-address/prefix-length
        — rpf-select
          — [no] core-mvpn service-id
            — group-prefix ip-address/mask [ip-address/mask...(upto 8 max)] [starg]
            — no group-prefix ip-address/mask
        — provider-tunnel
          — inclusive
            — 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
            — [no] wildcard-spmsi
          — selective
            — [no] auto-discovery-disable
            — 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] **pim-asm** {*grp-ip-address/mask*|*grp-ip-address netmask*}
- **pim-ssm** {*grp-ip-address/mask*|*grp-ip-address net-mask*}
- **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...*(up to 5 max)]}
- **no vrf-export**
- **vrf-import** {**unicast** | *policy-name* [*policy-name...*(up to 5 max)]}
- **no vrf-import**
- **vrf-target** {**unicast** | *ext-community* | **export unicast** | *ext-community* | **import unicast** | *ext-community*}
- **no vrf-target**
 - **export** {**unicast** | *ext-community*}
 - **import** {**unicast** | *ext-community*}

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
        — [no] description description-string
        — [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

```

Router Advertisement Commands

```

config
  — service
    — vprn
      — [no] router-advertisement
        — [no] dns-options
          — dns-servers ipv6-address
          — no dns-servers
          — rdnss-lifetime {seconds | infinite}
          — no rdnss-lifetime
        — [no] interface ip-int-name
          — current-hop-limit number
          — no current-hop-limit
          — [no] dns-options
            — dns-servers ipv6-address
            — no dns-servers
            — 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

```

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
        — arp-limit limit [log-only] [threshold percent]
        — no arp-limit
        — [no] arp-populate
        — arp-retry-timer timer-multiple
        — no arp-retry-timer
        — arp-timeout [seconds]
        — no arp-timeout
        — authentication-policy name
        — no authentication-policy
        — bfd transmit-interval [receive receive-interval] [multiplier multiplier][echo-receive echo-interval]
        — no bfd
        — cflowd-parameters
        — no cflowd-parameters
          — sampling {unicast | multicast} type {acl | interface} [direction {ingress-only | egress-only|both}]
          — no sampling {unicast | multicast}
        — delayed-enable seconds
        — no delayed-enable
        — description description-string
        — no description
        — 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}]
        — 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]
        — if-attribute
          — [no] admin-group group-name [group-name...(up to 5 max)]
          — srlg-group group-name [group-name...(up to 5 max)]
          — no srlg-group
        — ingress
          — policy-accounting <template-name>
          — no policy-accounting

```

- **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**] [**preferred**]
 - **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] dhcp6-server**
 - **icmp6**
 - **link-local-address** *ipv6-address* [**preferred**]
 - **no link-local-address**
 - **[no] local-proxy-nd**
 - **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...(up to 5 max)*]
 - **no proxy-nd-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**
- **[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*}
- **[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**

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]
        — no bfd
        — cflowd-parameters
        — no cflowd-parameters
          — sampling {unicast | multicast} type {acl | interface} [direction {ingress-only | egress-only|both}]
          — no sampling {unicast | multicast}
        — delayed-enable seconds
        — no delayed-enable
        — description 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]
        — 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]
          — [no] flowspec-ipv6
        — ip-mtu octets
        — no ip-mtu
        — lag lag-id[:encap-val]
        — no lag
        — [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] spi-load-balancing**
- **[no] teid-load-balancing**
- **mac** *ieee-address*
- **no mac**
- **[no] ntp-broadcast**
- **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**] [**igp-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**}

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]
          — arp arpId type type
          — no arp
          — accounting-policy acct-policy-id
          — no accounting-policy
          — bfd-template name
          — no bfd-template
          — [no] bfd-enable
          — egress
            — filter ip ip-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]
          — eth-cfm
            — [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 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} [cre-
              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
            — [no] squelch-ingress-levels [md-level [md-level...]]
          — [no] flowspec
          — [no] flowspec-ipv6
          — ingress

```

- **filter ip** *ip-filter-id*
- **filter ipv6** *ipv6-filter-id*}
- **no filter**
- **[no] flowspec**
- **[no] flowspec**
- **qos** *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**

Interface VRRP Commands

```

config
  — service
    — vprn service-id [customer customer-id]
    — no vprn service-id
      — interface ip-int-name
        — ipv6
          — vrrp virtual-router-id [owner]
          — no vrrp virtual-router-id
            — [no] backup ip-address
            — bfd-enable interface interface-name dst-ip ip-address
            — bfd-enable service-id interface interface-name dst-ip ip-address
            — no bfd-enable interface interface-name dst-ip ip-address
            — no bfd-enable service-id interface interface-name dst-ip ip-address
            — init-delay seconds
            — no init-delay
            — mac ieee-address
            — no mac
            — [no] master-int-inherit
            — message-interval {[seconds] [milliseconds milliseconds]}
            — no message-interval
            — [no] ping-reply
            — policy vrrp-policy-id
            — no policy
            — [no] preempt
            — priority priority
            — no priority
            — [no] shutdown
            — [no] ssh-reply
            — [no] standby-forwarding
            — [no] telnet-reply
            — [no] traceroute-reply
          — vrrp virtual-router-id [owner]
          — no vrrp virtual-router-id
            — authentication-key {authentication-key | hash-key} [hash | hash2]
            — no authentication-key
            — [no] backup ip-address
            — [no] bfd-enable [service-id] interface interface-name dst-ip ip-address
            — init-delay seconds
            — no init-delay
            — mac ieee-address
            — no mac
            — [no] master-int-inherit
            — message-interval {[seconds] [milliseconds milliseconds]}
            — no message-interval
            — [no] ping-reply
            — policy vrrp-policy-id

```

- **no policy**
- **[no] preempt**
- **priority** *priority*
- **no priority**
- **[no] shutdown**
- **[no] ssh-reply**
- **[no] standby-forwarding**
- **[no] telnet-reply**
- **[no] traceroute-reply**

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
          — arp aarpId type type
          — no arp
          — accounting-policy acct-policy-id
          — no accounting-policy [acct-policy-id]
          — anti-spoof {ip | mac | ip-mac}
          — no anti-spoof
          — calling-station-id calling-station-id
          — no calling-station-id
          — [no] bfd-enable
          — 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]
            — 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-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
                — high-prio-only percent
                — no high-prio-only
                — mbs {size-in-kbytes | default}
                — no mbs
                — 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**
 - **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**
 - **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]
 - **[no] flowspec**
 - **[no] flowspec-ipv6**
 - **host-lockout-policy** *policy-name*
 - **no host-lockout-policy**
 - **[no] host-shutdown**
 - **ingress**
 - **filter ip** *ip-filter-id*
 - **no filter** [*ip ip-filter-id*]
 - **[no] flowspec**
 - **[no] flowspec**
 - **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**
 - **avg-frame-overhead** *percentage*
 - **no avg-frame-overhead**
 - **cbs** *size-in-kbytes*
 - **no cbs**

- **high-prio-only** *percent*
- **no high-prio-only**
- **mbs** {*size-in-kbytes* | **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**
-
- **lag-link-map-profile** *lag-ink-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**
- **tod-suite** *tod-suite-name*
- **no tod-suite**
- **transit-policy** *ip-aasub-policy-id*
- **transit-policy** **prefix** *prefix-aasub-policy-id*
- **no transit-policy**

Oper Group Commands

```

config
  — service
    — vprn service-id
      — site name [create]
        — monitor-oper-group name
        — no monitor-oper-group name

```

Network Ingress Commands

```

config
  — service
    — vprn [inter-as-mvpn] [customer customer-id] [create]
    — no vprn service-id
      — network
        — ingress
          — qos network-policy-id fp-redirect-group queue-group-name
            instance instance-id
          — no qos

```

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]
        — best-path-selection
          — always-compare-med [zero|infinity]
          — always-compare-med strict-as {zero|infinity}
          — no always-compare-med
          — as-path-ignore [ipv4] [ipv6]
          — no as-path-ignore
          — [no] deterministic-med
        — [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]
        — no disable-communities
        — [no] disable-fast-external-failover
        — [no] eibgp-loadbalance
        — enable-bgp-vpn-backup [ipv4] [ipv6]
        — no enable-bgp-vpn-backup
        — [no] enable-peer-tracking
        — error-handling
          — [no] update-fault-tolerance
        — export policy-name [policy-name...(upto 5 max)]
        — no export
        — family [ipv4][ipv6] [mcast-ipv4] [flow-ipv4]
        — no family
        — flowspec-validate
        — [no] flowspec-validate
        — family
        — [no] graceful-restart
          — enable-notification
          — restart-time seconds
          — [no] stale-routes-time time
        — hold-time seconds [min seconds2]
        — no hold-time

```

- **[no] ibgp-multipath**
- **import** *policy-name* [*policy-name...*(up to 5 max)]
- **no import**
- **keepalive** *seconds*
- **no keepalive**
- **local-as** *as-number* [**private**]
- **no local-as**
- **local-preference** *local-preference*
- **no local-preference**
- **loop-detect** {**drop-peer** | **discard-route** | **ignore-loop**| **off**}
- **no loop-detect**
- **med-out** {**number** | **igp-cost**}
- **no med-out**
- **min-route-advertisement** *seconds*
- **no min-route-advertisement**
- **multihop** *ttl-value*
- **no multihop**
- **multipath** *max-paths* [**eibgp**]
- **no multipath**
- **next-hop-resolution**
 - **policy** *policy-name*
- **peer-tracking-policy** *policy-name*
- **preference** *preference*
- **no preference**
- **[no] rapid-withdrawal**
- **remove-private** [**limited**] [**skip-peer-as**]
- **no remove-private**
- **rib-management**
 - **ipv4**
 - **leak-import** *policy-name* [*policy-name .. (up to 15 max)*]
 - **no leak-import**
 - **ipv6**
 - **leak-import** *policy-name* [*policy-name .. (up to 15 max)*]
 - **no leak-import**
- **router-id** *ip-address*
- **no router-id**
- **[no] third-party-nexthop**
- **[no] shutdown**
- **[no] split-horizon**
- **[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-client-reflect**
- **disable-communities** [**standard**] [**extended**]
- **no disable-communities**
- [no] **disable-fast-external-failover**
- **ebgp-link-bandwidth** [**ipv4**] [**ipv6**]
- [no] **enable-peer-tracking**
- **error-handling**
 - [no] **update-fault-tolerance**
- **export** *policy-name* [*policy-name...(upto 5 max)*]
- **no export**
- **family** [**ipv4**][**ipv6**] [**mcast-ipv4**] [**flow-ipv6**] [**flow-ipv4**]
- **no family**
- **flowspec-validate**
- [no] **flowspec-validate**
- [no] **graceful-restart**
 - **enable-notification**
 - **restart-time** *seconds*
 - [no] **stale-routes-time** *time*
- **hold-time** *seconds* [**min** *seconds2*]
- **no hold-time**
- **import** *policy-name* [*policy-name...(upto 5 max)*]
- **no import**
- **keepalive** *seconds*
- **no keepalive**
- **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**}] [**post-import**]
- **no prefix-limit** *family*
- [no] **remove-private**
- [no] **shutdown**

- **[no] third-party-nexthop**
- **ttl-security** *min-ttl-value*
- **no ttl-security**
- **third-party-nexthop** {**internal** | **external**}
- **no third-party-nexthop**
- **[no] updated-error-handling**
- **[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**]
 - **no disable-communities**
 - **[no] disable-fast-external-failover**
 - **ebgp-link-bandwidth** [**ipv4**] [**ipv6**]
 - **[no] enable-peer-tracking**
 - **error-handling**
 - **[no] update-fault-tolerance**
 - **export** *policy-name* [*policy-name...*(upto 5 max)]
 - **no export**
 - **[no] graceful-restart**
 - **enable-notification**
 - **restart-time** *seconds*
 - **[no] stale-routes-time** *time*
 - **[no] stale-routes-time** *time*
 - **family** [**ipv4**][**ipv6**] [**mcast-ipv4**] [**flow-ipv6**] [**flow-ipv4**]
 - **no family**
 - **flowspec-validate**
 - **[no] flowspec-validate**
 - **hold-time** *seconds* [**min** *seconds2*]
 - **no hold-time**
 - **import** *policy-name* [*policy-name...*(upto 5max)]
 - **no import**
 - **family** [**ipv4**]
 - **keepalive** *seconds*
 - **no keepalive**
 - **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**}] [**post-import**]
- **no prefix-limit** *family*
- [**no**] **remove-private**
- [**no**] **shutdown**
- [**no**] **third-party-nexthop**
- **ttl-security** *min-ttl-value*
- **no ttl-security**
- **third-party-nexthop** {**internal** | **external**}
- **no third-party-nexthop**
- [**no**] **updated-error-handling**

ISIS Configuration Commands

```

config
  — service
    — vprn service-id [customer customer-id]
    — no vprn service-id
      — [no] isis [isis-instance]
        — [no] advertise-passive-only
        — [no] advertise-router-capability {area | as}
        — 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
        — [no] default-route-tag tag
        — export policy-name [policy-name...(upto 5 max)]
        — no export
        — export-limit number [log percentage]
        — no export-limit
        — [no] graceful-restart
          — [no] helper-disable
        — [no] hello-authentication
        — hello-padding {adaptive|loose|strict}
        — no hello-padding
        — ignore-attached-bit
        — no ignore-attached-bit
        — [no] ignore-lsp-errors
        — [no] iid-tlv-enable
        — [no] interface ip-int-name
          — bfd-enable {ipv4 | ipv6} [include-bfd-tlv]
          — no bfd-enable
          — csnp-interval seconds
          — no csnp-interval
          — 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
          — interface-type {broadcast|point-to-point}
          — no interface-type
          — [no] ipv4-multicast-disable
          — [no] ipv6-unicast-disable
          — level
            — 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**
- **ipv6-unicast-metric** *ipv6-metric*
- **no ipv6-unicast-metric**
- **ipv4-multicast-metric** *IPv4 multicast metric*
- **no ipv4-multicast-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**
- [no] **loopfree-alternate-exclude**
- **lsp-pacing-interval** *milli-seconds*
- **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**
- [no] **ipv4-routing**
- [no] **ipv6-routing**
- **level** *level*
 - [no] **advertise-router-capability**
 - [no] **auth-keychain** *name*
 - **authentication-key** *authentication-key*\hash-key [hash|hash2]
 - **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**
- **[no] external-preference** *external-preference*
- **[no] hello-authentication**
- **[no] loopfree-alternate-exclude**
- **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** [*1..8*]
 - **no oper-members**
 - **revert-members** [*1..8*]
 - **no revert-members**
- **[no] loopfree-alternate**
- **loopfree-alternate-exclude** *prefix-policy prefix-policy*
- **no loopfree-alternate-exclude**
- **lsp-lifetime** *seconds*
- **no lsp-lifetime**
- **lsp-mtu-size** *size*
- **no lsp-mtu-size**
- **lsp-refresh-interval** *seconds*
- **no lsp-refresh-interval**
- **lsp-wait** *lsp-wait [lsp-initial-wait [lsp-second-wait]]*
- **[no] multi-topology**
 - **[no] ipv4-multicast**
 - **[no] ipv6-unicast**
- **[no] multicast-import** [*ipv4*]
- **overload** [*timeout seconds*] [**max-metric**]
- **no overload**
- **overload-on-boot** [*timeout seconds*] [**max-metric**]
- **no overload-on-boot**
- **[no] poi-tlv-enable**
- **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**
- **[no] router-id** *router-id*
- **[no] rsvp-shortcut**
- **[no] shutdown**
- **[no] spf-wait** *spf-wait [spf-initial-wait [spf-second-wait]]*
- **[no] strict-adjacency-check**
- **summary-address** {*ip-prefix/mask* | *ip-prefix [netmask]*} *level* [**tag** *tag*]

- **no summary-address** {*ip-prefix/mask* | *ip-prefix* [*netmask*]}
- **suppress-attached-bit**
- **no suppress-attached-bit**
- **[no] unicast-import-disable**

OSPF Configuration Commands

```

config
  — service
    — vprn service-id [customer customer-id]
    — no vprn service-id
      — [no] ospf
      — ospf3 [instance-id] [router-id]
      — [no] ospf3 instance-id
        — advertise-router-capability { link | area | as }
        — no advertise-router-capability
        — [no] area area-id
          — area-range ip-prefix/mask [advertise | not-advertise]
          — no area-range ip-prefix/mask
          — [no] blackhole-aggregate
          — [no] interface ip-int-name [secondary]
            — [no] advertise-subnet
            — authentication bidirectional sa-name
            — authentication inbound sa-name outbound sa-name
            — no authentication
            — 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
            — [no] graceful-restart
              — helper-disable
            — hello-interval seconds
            — no hello-interval
            — interface-type {broadcast | point-to-point}
            — no interface-type
            — lfa-policy-map route-nh-template template-name
            — no lfa-policy-map
            — [no] loopfree-alternate-exclude
            — 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] passive
            — 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* [**advertise** | **not-advertise**]
 - **no area-range** *ip-prefix/mask*
 - **originate-default-route** [**type-7**]
 - **no originate-default-route**
 - **[no] redistribute-external**
 - **[no] summaries**
- **[no] sham-link** *ip-int-name ip-address*
 - **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
 - **no authentication-key**
 - **authentication-type** {**password** | **message-digest**}
 - **no authentication-type**
 - **dead-interval** *seconds*
 - **no dead-interval**
 - **hello-interval** *seconds*
 - **no hello-interval**
 - **message-digest-key** *key-id md5* [*key* | *hash-key*] [**hash** | **hash2**]
 - **no message-digest-key** *key-id*
 - **metric** *metric*
 - **no metric**
 - **retransmit-interval** *seconds*
 - **no retransmit-interval**
 - **[no] shutdown**
 - **transit-delay** *seconds*
 - **no transit-delay**
- **[no] stub**
 - **default-metric** *metric*
 - **no default-metric**
 - **[no] summaries**
- **[no] virtual-link** *router-id transit-area area-id*
 - **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
 - **no authentication-key**
 - **authentication-type** {**password** | **message-digest**}
 - **no authentication-type**
 - **dead-interval** *seconds*
 - **no dead-interval**
 - **hello-interval** *seconds*
 - **no hello-interval**
 - **message-digest-key** *key-id md5* [*key* | *hash-key*] [**hash** | **hash2**]
 - **no message-digest-key** *key-id*
 - **retransmit-interval** *seconds*
 - **no retransmit-interval**
 - **[no] shutdown**
 - **transit-delay** *seconds*
 - **no transit-delay**

- **[no] compatible-rfc1583**
- **export** *policy-name* [*policy-name...*(up to 5 max)]
- **no export**
- **external-db-overflow** *limit seconds*
- **no external-db-overflow**
- **external-preference** *preference*
- **no external-preference**
- **[no] graceful-restart**
 - **[no] helper-disable**
- **[no] ignore-dn-bit**
- **[no] ignore-dn-bit**
- **[no] loopfree-alternate**
- **loopfree-alternate-exclude** *prefix-policy prefix-policy* [*prefix-policy...* up to 5]
- **no loopfree-alternate-exclude**
- **[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**
- **[no] shutdown**
- **[no] super-backbone**
- **[no] suppress-dn-bit**
- **timers**
 - **[no] lsa-arrival** *lsa-arrival-time*
 - **[no] lsa-generate** *max-lsa-wait* [*lsa-initial-wait* [*lsa-second-wait*]]
 - **[no] spf-wait** *max-spf-wait* [*spf-initial-wait* [*spf-second-wait*]]
- **[no] unicast-import-disable**
- **vpn-domain** *id* {**0005** | **0105** | **0205** | **8005**}
- **no vpn-domain**
- **vpn-tag** *vpn-tag*
- **no vpn-tag**

PIM Configuration Commands

```

config
  — service
    — vprn
      — [no] pim
        — apply-to {all | none}
        — [no] grt-extranet
          — group-prefix ip-address/mask [ip-address/mask...(up to 8 max)] [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
          — mcaac
            — mc-constraints
              — level level-id bw bandwidth
              — no level
              — 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
          — 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**
- **[no] mc-ecmp-balance-hold**
- **[no] non-dr-attract-traffic**
- **rp**
 - **auto-rp-discovery**
 - **[no] anycast** *rp-ip-address*
 - **[no] rp-set-peer** *ip-address*
 - **[no] auto-rp-discovery**
 - **bootstrap-export** *policy-name* [*.. policy-name...up to five*]
 - **no bootstrap-export**
 - **bootstrap-import** *policy-name* [*.. policy-name...up to five*]
 - **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**
 - **anycast** *ipv6-address*
 - **[no] rp-set-peer** *ipv6-address*
 - **bsr-candidate** *ipv6-address*
 - **address** *ipv6-address*
 - **[no] address**
 - **hash-mask-length** *hash-mask-length*
 - **[no] hash-mask-length**
 - **priority** *bootstrap-priority*
 - **no priority**
 - **[no] shutdown**
 - **[no] embedded-rp**
 - **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**
- [no] **rpf-table** {rtable-m | rtable-u | both}
- [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}

C-Multicast Listener Discovery (MLD) Commands

- config
 - service
 - [no] vprn
 - [no] **mld**
 - [no] **interface** ip-int-name
 - [no] **disable-router-alert-check**
 - **import** policy-name
 - **no import**
 - **max-groups** value
 - **no max-groups**
 - **mcac**
 - **mc-constraints**
 - **level** level-id **bw** bandwidth
 - **no level**
 - **number-down** number-lag-port-down **level** level-id
 - **no number-down**
 - [no] **shutdown**
 - [no] **use-lag-port-weight**
 - **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] 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*

Show Commands

```

show
  — service
    — egress-label start-label [end-label]
    — ingress-label start-label [[end-label]
    — id service-id
      — all
      — authentication
        — statistics [policy name] [sap sap-id]
      — arp [ip-address] | [mac ieee-address] | [sap port-id:encap] | [interface ip-int-name]
      — 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]
      — base
      — sessions summary
      — sessions [detail] wholesaler service-id
      — interface [ip-address | ip-int-name] [detail]
      — ptp
      — retailers
      — sap [sap-id] [detail]
      — sdp [sdp-id | far-end ip-address] [detail]
    — sap-using [sap sap-id]
    — sap-using interface [ip-address | ip-int-name]
    — sap-using [ingress | egress] atm-td-profile id-profile-id
    — sap-using [ingress | egress] filter filter-id
    — sap-using [ingress | egress] qos-policy qos-policy-id
    — sap-using authentication-policy policy-name
    — sap-using [sdp-id | far-end ip-address] [detail | keep-alive-history]
    — sap-using [sdp-id[:vc-id]
    — service-using [vprn] [sdp sdp-id] [customer customer-id]
    — sdp
      — 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
    — pw-port
      — pw-port [pw-port-id] [detail]
      — pw-port sdp [sdp-id]
      — pw-port sdp none

show
  — service 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 [regexp]] [ascending|descending]
      — snmp-trap-group [log-id]
      — syslog [syslog-id]
  
```

```

show
— router [vprn-service-id]
  — aggregate [active]
  — arp [ip-address | ip-int-name | mac ieee-mac-address] [sdp sdp-id:vc-id] [summary]
  — bgp
    — damping [ip-prefix/mask | ip-address] [detail]
    — damping [damp-type] [detail]
    — group [name] [detail]
    — neighbor [ip-address [[family family] filter1]]
    — neighbor [as-number [[family family] filter2]]
    — paths
    — routes [family family] [prefix [detail | longer]]
    — routes [family family] [prefix [hunt | brief]]
    — routes [family family] [community comm-id]
    — routes [family family] [aspath-regex reg-ex1]
    — routes [family] [ipv6-prefix[/prefix-length] [detail | longer]][[hunt [brief]]]
    — summary [all]
  — ecmp
  — interface [{[ip-address | ip-int-name] [detail]} | summary | exclude-services]
  — ldp
    — bindings active
  — mvpn
  — route-table [ip-address[/mask] [longer | best]] | [protocol protocol] | [summary]
  — service-prefix
  — static-arp [ip-address | ip-int-name | mac ieee-mac-address]
  — static-route [ip-prefix /mask] | [preference preference] | [next-hop ip-address] [detail]
  — tunnel-table [ip-address[/mask] [protocol protocol | sdp sdp-id]
  — tunnel-table [summary]

```

Clear Commands

```

clear
  — router
    — arp-host
    — arp-host { mac ieee-address | sap sap-id | ip-address ip-address[/mask] }
    — arp-host [port port-id] [inter-dest-id intermediate-destination-id | no-inter-dest-id]
    — arp-host statistics [sap sap-id | interface interface-name]
    — bgp
      — damping [{prefix/mask [neighbor ip-address]} | {group name}]
      — flap-statistics [[ip-prefix/mask] [neighbor ip-address]] | [group group-name] |
        [regex reg-exp] | [policy policy-name]
      — neighbor {ip-address | as as-number | external | all} [soft | soft-inbound | statis-
        tics]
      — protocol
    — dhcp
      — statistics [interface ip-int-name | ip-address]
    — interface [ip-int-name | ip-address] [icmp]
    — rip

clear
  — service
    — id service-id
      — dhcp
        — lease-state
        — lease-state ip-address ip-address
        — lease-state mac ieee-address
        — lease-state sap sap-id
        — lease-state sdp sdp-id:vc-id
      — spoke-sdp sdp-id:vc-id ingress-vc-label
    — statistics
      — sap sap-id {all | counters | stp}
      — sdp sdp-id keep-alive
      — id service-id
        — counters
        — spoke-sdp sdp-id:vc-id {all | counters | stp}
        — spoke-sdp
  
```

Debug Commands

```

debug
  — service
    — idservice-id
      — [no] arp-host
      — [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 {config-change | oper-status-change}
      — [no] sdp sdp-id:vc-id
        — event-type {config-change | oper-status-change}
      — stp
        — [no] all-events
        — [no] bpdu
        — [no] exception
        — [no] exception
        — [no] fsm-state-changes
        — [no] fsm-timers
        — [no] port-role
        — [no] port-state
        — [no] sap sap-id
        — [no] sdp sdp-id:vc-id

debug
  — router [router-instance]
    — igmp
      — [no] interface [ip-int-name | ip-address]
      — [no] mcs [ip-int-name]
      — [no] misc
      — [no] packet [query|v1-report|v2-report|v3-report|v2-leave] [ip-int-name|ip-address]

```

VPRN Service Configuration Commands

Generic Commands

shutdown

Syntax [no] shutdown

Context config>service>vprn
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>l2tp

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

Generic Commands

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

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

Default No description associated with the configuration context.

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

Global Commands

vprn

Syntax	vprn <i>service-id</i> [customer <i>customer-id</i>] [create] no vprn <i>service-id</i>						
Context	config>service						
Description	<p>This command creates or edits a Virtual Private Routed Network (VPRN) service instance.</p> <p>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.</p> <p>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.</p> <p>IP interfaces defined within the context of an VPRN service ID must have a SAP created as the access point to the subscriber network.</p> <p>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.</p> <p>When a service is created, the use of the customer <i>customer-id</i> is optional to navigate into the service configuration context. If attempting to edit a service with the incorrect <i>customer-id</i> results in an error.</p> <p>Multiple VPRN services are created to separate customer-owned IP interfaces. More than one VPRN service can be created for a single customer ID. More than one IP interface can be created within a single VPRN service ID. All IP interfaces created within an VPRN service ID belongs to the same customer.</p> <p>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.</p>						
Default	None — No VPRN service instances exist until they are explicitly created.						
Parameters	<p><i>service-id</i> — The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The <i>service-id</i> must be the same number used for every on which this service is defined.</p> <table> <tr> <td>Values</td> <td><i>service-id</i>:</td> <td>1 — 2147483648</td> </tr> <tr> <td></td> <td><i>svc-name</i>:</td> <td>64 characters maximum</td> </tr> </table> <p>customer <i>customer-id</i> — 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.</p>	Values	<i>service-id</i> :	1 — 2147483648		<i>svc-name</i> :	64 characters maximum
Values	<i>service-id</i> :	1 — 2147483648					
	<i>svc-name</i> :	64 characters maximum					

Values 1 — 2147483647

aggregate

Syntax	<p>aggregate <i>ip-prefix/ip-prefix-length</i> [summary-only] [as-set] [aggregator <i>as-number:ip-address</i>] [black-hole] [community <i>comm-id</i>] [description <i>description</i>]</p> <p>aggregate <i>ip-prefix/ip-prefix-length</i> [summary-only] [as-set] [aggregator <i>as-number:ip-address</i>] [community <i>comm-id</i>] [indirect <i>ip-address</i>] [description <i>description</i>]</p> <p>no aggregate <i>ip-prefix/ip-prefix-length</i></p>				
Context	config>service>vprn				
Description	<p>This command creates an aggregate route.</p> <p>Use this command to automatically install an aggregate in the routing table when there are one or more component routes. A component route is any route used for forwarding that is a more-specific match of the aggregate.</p> <p>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.</p> <p>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.</p> <p>A standard 4-byte BGP community may be associated with an aggregate route in order to facilitate route policy matching.</p> <p>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.</p> <p>The no form of the command removes the aggregate.</p>				
Default	No aggregate routes are defined.				
Parameters	<p><i>ip-prefix</i> — The destination address of the aggregate route in dotted decimal notation.</p> <table border="0" style="margin-left: 2em;"> <tr> <td style="vertical-align: top;">Values</td> <td> <p>ipv4-prefix a.b.c.d (host bits must be 0)</p> <p>ipv4-prefix-length 0 — 32</p> <p>ipv6-prefix x:x:x:x:x:x:x (eight 16-bit pieces)</p> <p> x:x:x:x:x:d.d.d.d</p> <p> x: [0 — FFFF]H</p> <p> d: [0 — 255]D</p> <p>ipv6-prefix-length 0 — 128</p> </td> </tr> </table> <p>The mask associated with the network address expressed as a mask length.</p> <table border="0" style="margin-left: 2em;"> <tr> <td style="vertical-align: top;">Values</td> <td>0 — 32</td> </tr> </table> <p>summary-only — This optional parameter suppresses advertisement of more specific component routes for the aggregate.</p>	Values	<p>ipv4-prefix a.b.c.d (host bits must be 0)</p> <p>ipv4-prefix-length 0 — 32</p> <p>ipv6-prefix x:x:x:x:x:x:x (eight 16-bit pieces)</p> <p> x:x:x:x:x:d.d.d.d</p> <p> x: [0 — FFFF]H</p> <p> d: [0 — 255]D</p> <p>ipv6-prefix-length 0 — 128</p>	Values	0 — 32
Values	<p>ipv4-prefix a.b.c.d (host bits must be 0)</p> <p>ipv4-prefix-length 0 — 32</p> <p>ipv6-prefix x:x:x:x:x:x:x (eight 16-bit pieces)</p> <p> x:x:x:x:x:d.d.d.d</p> <p> x: [0 — FFFF]H</p> <p> d: [0 — 255]D</p> <p>ipv6-prefix-length 0 — 128</p>				
Values	0 — 32				

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 *comm-id* — 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.

Values	comm-id	asn:comm-val well-known-comm
	asn	0 — 65535
	comm-val	0 — 65535
	well-known-comm	no-advertise, no-export, no-export-subconfed

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.

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
		x:x:x:x:x:d.d.d.d
		x: [0 — FFFF]H
		d: [0 — 255]D

description *description-text* — Specifies a text description stored in the configuration file for a configuration context.

allow-export-bgp-vprn

Syntax	allow-export-bgp-vprn no allow-export-bgp-vprn
Context	config>service>vprn
Description	This command causes the vrf-export and vrf-target functions of the VPRN to include BGP-VPN routes installed in the VPRN route table. For split-horizon reasons, these routes are normally not re-advertisable as VPN-IP routes. When a BGP-VPN route is re-exported, the route-distinguisher and label values are rewritten per the configuration of the re-exporting VPRN. Note — Appropriate measures should be taken to ensure that routing updates do not loop back to the source when this command is used, otherwise the routes could become unstable.
Default	no allow-export-bgp-vprn

auto-bind-tunnel

Syntax	auto-bind-tunnel
Context	config>service>vprn
Description	<p>This command enables the context to configure automatic binding of a VPRN service using tunnels to MP-BGP peers.</p> <p>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.</p> <p>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.</p> <p>The following tunnel types are supported in a VPRN context in order of preference: RSVP, LDP, Segment Routing (SR), and GRE. The BGP tunnel type is not explicitly configured and is thus implicit. It is always preferred over any other tunnel type enabled in the auto-bind-tunnel context.</p> <p>The ldp value instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.</p> <p>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 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.</p> <p>When the sr-isis (sr-ospf) value is enabled, a SR tunnel to the BGP next-hop is selected in the TTM from the lowest numbered ISIS (OSPF) instance.</p> <p>The user must set resolution to filter to activate the list of tunnel-types configured under resolution-filter.</p> <p>When an explicit SDP to a BGP next-hop is configured in a VPRN service (configure>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.</p>

resolution

Syntax	resolution {any filter disabled}
Context	config>service>vprn>auto-bind-tunnel
Description	<p>This command configures the resolution mode in the automatic binding of a VPRN service to tunnels to MP-BGP peers.</p>

- Parameters**
- any** — Enables the binding to any supported tunnel type in VPRN context following TTM preference.
 - filter** — Enables the binding to the subset of tunnel types configured under resolution-filter.
 - disabled** — Disables the automatic binding of a VPRN service to tunnels to MP-BGP peers.

resolution-filter

- Syntax** **resolution-filter**
- Context** config>service>vprn>auto-bind-tunnel
- Description** This command configures the subset of tunnel types which 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, LDP, Segment Routing (SR), and GRE. The BGP tunnel type is not explicitly configured and is thus implicit. It is always preferred over any other tunnel type enabled in the **auto-bind-tunnel** context.
- Parameters**
- gre** — Selects the GRE tunnel type.
 - ldp** — Selects the LDP tunnel type.
 - rsvp** — Selects the RSVP-TE tunnel type.
 - sr-isis** — Selects the Segment Routing (SR) tunnel type programmed by an IS-IS instance in TTM.

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** *as-number* — Specifies the AS number for the VPRN service.
- Values** 1 — 4294967295

backup-path

- Syntax** [**no**] **backup-path** [**ipv4**] [**ipv6**]
- Context** config>router

config>service>vprn

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. In some cases prefix independent convergence may require use of FP2 or later IOMs/IMMs/XMAs.

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 IPv4 prefixes

ipv6 — Enables the use of a backup path for BGP-learned IPv6 prefixes

carrier-carrier-vpn

Syntax	[no] carrier-carrier-vpn
Context	config>service>vpn
Description	<p>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.</p> <p>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.</p> <p>The no form of the command removes the Carrier Supporting Carrier capability from a VPRN.</p>
Default	no carrier-carrier-vpn

confederation

Syntax	confederation <i>confed-as-num</i> members <i>as-number</i> [<i>as-number...</i> (up to 15 max)] no confederation <i>confed-as-num</i> members <i>as-number</i> [<i>as-number...</i> (up to 15 max)] no confederation
Context	config>service>vpn
Description	<p>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.</p> <p>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).</p> <p>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.</p>
Default	No confederations are defined.
Parameters	<i>confed-as-num</i> — The confederation AS number defined as a decimal value. Values 1 — 4294967295 <i>members 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 — 4294967295

dns

Syntax	[no] dns
Context	config>service>vpn

Global Commands

Description This command enables 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** *ipv4-address*
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 *ipv4-address* 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.

Default none

Parameters *ipv4-address* — Specifies the IPv4 address of the default secondary DNS server.

Values *ipv4-address* - a.b.c.d

ipv6-source-address

Syntax **ipv6-source-address** *ipv6-address*
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 *ipv6-address* 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.

Default none

Parameters *ipv4-address* — Specifies the IPv6 address of the default secondary DNS server.

Values *ipv4-address* - 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	<i>ip-address</i> — 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:x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface - 32 chars max, for link local addresses.

secondary-dns

Syntax	secondary-dns <i>ip-address</i> 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	<i>ip-address</i> — 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:x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface - 32 chars max, for link local addresses

tertiary-dns

Syntax	tertiary-dns <i>ip-address</i> no tertiary-dns												
Context	config>service>vprn>dns												
Description	<p>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.</p> <p>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.</p> <p>The no form of the command removes the tertiary DNS server from the configuration.</p>												
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	<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[-interface]</td> </tr> <tr> <td></td> <td>x:x:x:x:x:d.d.d.d[-interface]</td> </tr> <tr> <td></td> <td>x: [0..FFFF]H</td> </tr> <tr> <td></td> <td>d: [0..255]D</td> </tr> <tr> <td></td> <td>interface - 32 chars max, for link local addresses</td> </tr> </table>	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..FFFF]H		d: [0..255]D		interface - 32 chars max, for link local addresses
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..FFFF]H												
	d: [0..255]D												
	interface - 32 chars max, for link local addresses												

ecmp

Syntax	ecmp <i>max-ecmp-routes</i> no ecmp
Context	config>service>vprn
Description	<p>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.</p> <p>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.</p> <p>ECMP can only be used for routes learned with the same preference and same protocol. See the discussion on preferences in the static-route command.</p> <p>When more ECMP routes are available at the best preference than configured by the <i>max-ecmp-routes</i> parameter, then the lowest next-hop IP address algorithm is used to select the number of routes configured.</p> <p>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.</p>
Default	no ecmp
Parameters	<i>max-ecmp-routes</i> — Specifies the maximum number of routes for path sharing.

Values 0—32

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.

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

export-grt

Syntax	export-grt <i>policy-name</i> [<i>policy-name</i> ...(up to 5 max)] no export-grt
Context	config>service>vprn>grt-lookup

Global Commands

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

export-limit

Syntax **export-limit** *num-routes*
no export-limit

Context config>service>vprn>grt-lookup
config>service>vprn>ospf
config>service>vprn>ospf3

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 *num-routes* — Specifies maximum number of routes that can be exported.

Values 0 — 1000

export-v6-limit

Syntax **export-v6-limit** *num-routes*
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 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 *num-routes* — Specifies maximum number of routes that can be exported.

Values 0 — 1000

allow-local-management

Syntax	[no] allow-local-management
Context	config>service>vprn>grt-lookup>enable-grt
Description	When enabled, both IPv4 and IPv6 base interfaces shall respond to leaked traffic from the VPRN.

static-route

Syntax	static-route <i>{ip-prefix/prefix-length ip-prefix netmask}</i> [preference preference] [metric metric] [enable disable] grt no static-route																				
Context	config>service>vprn>grt-lookup>enable-grt																				
Description	This command is a simplified version of the traditional static-route command pointing to the base routing instance. This instructs the route lookup function to look only in the GRT for a route matching destination static route and not look up the route in the local VPRN. The GRT keyword is a required parameter. The no form causes the feature into the default mode of primary lookup for all routes in the local VPRN and failing a match in the local VPRN, the lookup result in the GRT will be used.																				
Parameters	<i>ip-prefix/prefix-length</i> — Specifies the IPv4 prefix and prefix length. <table> <tr> <td>Values</td> <td>ip-prefix: a.b.c.d (host bits must be 0)</td> </tr> <tr> <td></td> <td>ipv4-prefix-length: [0..32]</td> </tr> <tr> <td></td> <td>ipv6-prefix 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> <tr> <td></td> <td>ipv6-prefix-length 0 — 128</td> </tr> </table> <i>netmask</i> — Specifies the netmask. <table> <tr> <td>Values</td> <td>a.b.c.d (network bits all 1 and host bits all 0)</td> </tr> </table> <i>preference</i> — Specifies the preference. <table> <tr> <td>Values</td> <td>1..255</td> </tr> </table> <i>metric</i> — Specifies the metric. <table> <tr> <td>Values</td> <td>0..65535</td> </tr> </table> enable disable — Keyword; specifies the state of the static-route. grt — Keyword; Global Route Table lookup.	Values	ip-prefix: a.b.c.d (host bits must be 0)		ipv4-prefix-length: [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-length 0 — 128	Values	a.b.c.d (network bits all 1 and host bits all 0)	Values	1..255	Values	0..65535
Values	ip-prefix: a.b.c.d (host bits must be 0)																				
	ipv4-prefix-length: [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-length 0 — 128																				
Values	a.b.c.d (network bits all 1 and host bits all 0)																				
Values	1..255																				
Values	0..65535																				

IGMP Commands

igmp

Syntax	[no] igmp
Context	config>service>vprn
Description	This command enables 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.
Default	none
Parameters	<i>ip-int-name</i> — Specifies the name of the IP interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. fwd-service <i>service-id</i> — 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 — 2147483650, svc-name up to 64 char maximum
Default	none

disable-router-alert-check

Syntax	[no] disable-router-alert-check
Context	config>service>vprn>igmp>gr-if config>service>vprn>igmp>if
Description	This command enables the IGMP router alert check option. The no form of the command disables the router alert check.

import

Syntax	import <i>policy-name</i> no import
Context	config>service>vprn>igmp>gr-if config>service>vprn>igmp>if
Description	This command specifies the policy that is to be applied on this interface.
Parameters	<i>policy-name</i> — Specify the policy to filter IGMP packets.

max-groups

Syntax	max-groups <i>value</i> no max-groups
Context	config>service>vprn>igmp>gr-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	<i>value</i> — Specifies the maximum number of groups for this interface. Values 1 — 16000

max-sources

Syntax	max-sources [1..1000] no max-sources
Context	config>service>vprn>igmp>gr-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	<i>sources</i> — Specifies the maximum number of sources for this interface. Values 1 — 1000

max-grp-sources

Syntax	max-grp-sources [1..32000] no max-grp-sources
Context	config>service>vprn>igmp>gr-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 — 32000 — Specifies the maximum number of group source. Values 1 — 32000

mcac

Syntax	mcac
Context	config>service>vprn>igmp>gr-if
Description	This command enables the context to configure multicast CAC parameters.

mc-constraints

Syntax	mc-constraints
Context	config>service>vprn>igmp>gr-if
Description	This command configures multicast CAC constraints.

policy

Syntax	policy <i>policy-name</i> no policy
Context	config>service>vprn>igmp>gr-if
Description	This command references the global channel bandwidth definition policy that is used for (H)mcac and HQoS Adjust. HQoS Adjustment is supported only with redirection enabled. In other words, the policy from the redirected interface is used for HQoS Adjustment.

Hierarchical mcac (Hmccac) is supported only with redirection enabled. In Hmccac, 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 Hmccac case the channel definition policy must be referenced under the redirected interface level.

Parameters	<i>policy-name</i> — Specifies the name of the global mcac channel definition policy defined under the hierarchy configure>router>mcac>policy .
Default	No policy is referenced.

unconstrained-bw

Syntax	unconstrained-bw <i>bandwidth</i> mandatory-bw <i>mandatory-bw</i> no unconstrained-bw
Context	config>service>vprn>igmp>gr-if
Description	This command configures unconstrained-bw for multicast cac policy on this interface. The no form of the command
Parameters	<i>bandwidth</i> — Specifies the bandwidth assigned for interface's multicast cac policy traffic in kilo-bits per second (kbps). Values 0 — 2147483647 mandatory-bw <i>mandatory-bw</i> — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilo-bits per second (kbps). 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 — 2147483647

query-src-ip

Syntax	query-src-ip <i>ip-address</i> no query-src-ip
Context	config>service>vprn>igmp>gr-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.

Global Commands

Default none

Parameters *ip-address* — Sets the source IPv4 address for all subscriber's IGMP queries.

sub-hosts-only

Syntax **[no] sub-hosts-only**

Context config>service>vprn>igmp>gr-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>gr-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** *version*
no version

Context config>service>vprn>igmp>gr-if

Description This command configures the version of IGMP.
The **no** form of the command

Parameters *version* — Specifies the IGMP version.

Values 1, 2 or 3

grp-if-query-src-ip

Syntax **grp-if-query-src-ip** *ip-address*
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.

Default none

interface

Syntax **interface** *ip-int-name*
no interface

Context config>service>vprn>igmp

Description This command enables the context to configure IGMP interface parameters.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Values 1 — 32 characters maximum

import

Syntax **import** *policy-name*
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 *policy-name* — 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, etc.), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

max-groups

Syntax **max-groups** *value*
no max-groups

Context config>service>vprn>igmp>if

Description This command specifies the maximum number of groups for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed.

Default 0, no limit to the number of groups.

Global Commands

Parameters *value* — Specifies the maximum number of groups for this interface.
Values 1 — 16000

mcac

Syntax **mcac**
Context config>service>vprn>if
config>service>vprn>pim>if
Description This command configures multicast CAC policy and constraints for this interface.
Default none

mc-constraints

Syntax **mc-constraints**
Context config>service>vprn>igmp>if>mcac
config>service>vprn>pim>if>mcac
Description This command enables the context to configure multicast CAC constraints.
Default none

level

Syntax **level** *level-id* **bw** *bandwidth*
no level *level-id*
Context config>service>vprn>igmp>if>mcac
config>service>vprn>pim>if>mcac
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 — 8
bandwidth — Specifies the bandwidth in kilobits per second (kbps) for the level.
Values 1 — 2147483647

number-down

Syntax	number-down <i>number-lag-port-down</i> level <i>level-id</i> no number-down
Context	config>service>vprn>igmp>if>mcac>mc-constraints config>service>vprn>pim>if>mcac>mc-constraints
Description	This command configures the number of ports down and level for interface's multicast CAC policy.
Default	not enabled
Parameters	<i>number-lag-port-down</i> — 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 — 64 (for 64-link LAG) 1 — 32 (for other LAGs) <i>level-id</i> — Specifies an entry for the multicast CAC policy constraint level configured on this system. Values 1 — 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>if>mcac config>service>vprn>pim>if>mcac
Description	This command configures the mulitcast 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, etc.), the entire string must be enclosed within double quotes.

unconstrained-bw

Syntax	unconstrained-bw <i>bandwidth</i> mandatory-bw <i>mandatory-bw</i> no unconstrained-bw
Context	config>service>vprn>igmp>if>mcac config>service>vprn>pim>if>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	<p><i>bandwidth</i> — The bandwidth assigned for interface's MCAC policy traffic, in kilo-bits per second (kbps).</p> <p>Values 0 — 2147483647</p> <p>mandatory-bw <i>mandatory-bw</i> — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilo-bits per second (kbps).</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>Values 0 — 2147483647</p>

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.
Default	none

group

Syntax	[no] group <i>grp-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.

Default none

Parameters *grp-ip-address* — Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group. The address must be in dotted decimal notation

source

Syntax **source**

Context config>service>vprn>igmp>if>static>group

Description This command specifies a 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 source(s) 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.

Default none

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.

Default none

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 <i>version</i> no version
Context	config>service>vprn>igmp>if
Description	<p>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.</p> <p>For IGMPv3, note that 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.</p>
Default	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	125
Parameters	<i>seconds</i> — The time frequency, in seconds, that the router transmits general host-query messages. Values 2 — 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	1
Parameters	<i>seconds</i> — Specifies the frequency, in seconds, at which query messages are sent. Values 1 — 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	10
Parameters	<i>seconds</i> — Specifies the the length of time to wait to receive a response to the host-query message from the host. Values 1 — 1023

Global Commands

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	2
Parameters	<i>robust-count</i> — Specifies the robust count value. Values 2 — 10

ssm-translate

Syntax	igmp
Context	config>service>vprn>igmp config>service>vprn>igmp>if
Description	This command enables 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. <i>end</i> — 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 <i>start</i> 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** *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.

igmp-host-tracking

Syntax **igmp-host-tracking**

Context config>service>vprn
config>service>vprn>sap

Description This command enables the context to configure IGMP host tracking parameters.

expiry-time

Syntax **expiry-time** *expiry-time*
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 *expiry-time* — Specifies the time, in seconds, that this system continues to track an inactive host.

Values 1 — 65535

import

Syntax **import** *policy-name*
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 *policy-name* — 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>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 — 196607

maximum-routes

Syntax	maximum-routes <i>number</i> [log-only] [threshold <i>percentage</i>] 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. Note that local , host , static and aggregate routes are not counted. Note that 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 — 2147483647
	log-only — This parameter specifies that if the maximum limit is reached, only log the event. log-only does not disable the learning of new routes.
	threshold <i>percentage</i> — 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 — 100

multicast-info-policy

Syntax **multicast-info-policy** *policy-name*
no multicast-info-policy

Context config>service>vprn

Description This command configures multicast information policy.

Parameters *policy-name* — Specifies the policy name.
Values 32 chars max

mc-maximum-routes

Syntax **mc-maximum-routes** *number* [**log-only**] [**threshold** *threshold*]

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 *number* — Specifies the maximum number of routes to be held in a VRF context.
Values 1 — 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 *threshold* — The percentage at which a warning log message and SNMP trap should be sent.
Values 0 — 100
Default 10

network

Syntax **network**

Context config>service>vprn

Description This command enables the context to configure network parameters for the VPRN service.

ingress

Syntax	ingress
Context	config>service>vprn>network
Description	This command enables 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/ipv6-filter-id</i> — Specifies an existing IP/IPv6 filter policy of a scope template. Values [1..65535] <i>name</i> <i>name</i> : 64 characters maximum

ptp

Syntax	[no] ptp
Context	config>service>vprn
Description	This command enables the context to configure PTP parameters for the VPRN service.

peer-limit

Syntax	peer-limit <i>limit</i> no peer-limit
Context	configure>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 — 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 configure>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.
Default	none
Parameters	<i>a.b.c.d</i> — The IP address of the remote peer.
Values	ipv4-address a.b.c.d
Description	

reassembly-group

Syntax	reassembly-group <i>nat-group-id</i> no reassembly-group
Context	configure>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 route-distinguisher
Parameters	<i>nat-group-id</i> — Nat-group id. The nat-group contains up to 10 active ISAs.

asn:number — 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	<p>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.</p> <p>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.</p>
Default	no route-distinguisher
Parameters	<p>The route distinguisher is a 6-byte value that can be specified in one of the following formats:</p> <p><i>ip-address:number</i> — Specifies the IP address in dotted decimal notation. The assigned number must not be greater than 65535.</p> <p><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.</p> <p>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.</p>

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	<p>This command sets the router ID for a specific VPRN context.</p> <p>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:</p> <ul style="list-style-type: none">• 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. <p>If neither the router ID nor system interface are defined, the router ID from the base router context is inherited.</p>

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 node 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 *ip-address* — The IP address must be given in dotted decimal notation.

service-name

Syntax **service-name** *service-name*
no service-name

Context config>service>vprn

Description This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a given service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the and platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a given service once it is initially created.

Parameters *service-name* — Specifies a unique service name to identify the service. Service names may not begin with an integer (0-9).

sgt-qos

Syntax	sgt-qos
Context	config>service>vprn
Description	This command enables the context to configure DSCP/Dot1p re-marking for self-generated traffic.

application

Syntax	application <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 re-marking 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 IEEE802.1 dot1P and LSP EXP bits.
- The Dot1P and the LSP EXP bits are set by the egress complex for all packets based on the signaled FC. This includes ARP and IS-IS packets that, due to their nature, do not carry DSCP bits.
- The DSCP value in the egress IP header will be as configured in this command. The egress QoS policy will not overwrite this value.

Only one DSCP name/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 ldp, rsvp, bgp, rip, msdp, pim, ptp, ospf, igmp, mld, telnet, tftp, ftp, ssh, snmp, snmp-notification, syslog, icmp, traceroute, tacplus, dns, ntp, radius, cflowd, dhcp, bootp, ndis, vrrp, srrp
	<i>dscp-value</i> — Specifies a value when this packet egresses the respective egress policy should provide the mapping for the DSCP value to either LSP-EXP bits or IEEE 802.1p (Dot1P) bits as appropriate otherwise the default mapping applies.
	Values 0 — 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

dot1p-priority — Specifies the Dot1P priority.

Values 0 — 7

dot1p-app-name — Specifies the Dot1P application name.

Values arp, isis

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	<p>This command creates a mapping between the DiffServ Code Point (DSCP) of the self generated traffic and the forwarding class.</p> <p>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.</p> <p>All DSCP names that defines a DSCP value must be explicitly defined.</p> <p>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.</p>
Default	none
Parameters	<p><i>dscp-name</i> — The name of the DiffServ code point to be associated with the forwarding class. DiffServ code point can only be specified by its name and only an existing DiffServ code point can be specified. The software provides names for the well known code points.</p> <p>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</p> <p>fc <i>fc-name</i> — Specifies the forwarding class name. All packets with DSCP value or MPLS EXP bits that is not defined will be placed in this forwarding class.</p> <p>Default None, the fc name must be specified</p> <p>Values be, l2, af, l1, h2, ef, h1, nc</p>

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: <ul style="list-style-type: none"> • 7750 SR-12/SR-7/SR-c12 and 7450 ESS-12/ESS-7/ESS-6 platforms: protocol sets overload if one of the SF/CPMs fails • 7950 XRS and 7750 SR-12e platforms: protocol sets overload if two SFMs fail 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> — This parameter specifies the delay between detecting SFM failures and setting overload. <table> <tr> <td>Values</td> <td>1 — 600 seconds</td> </tr> <tr> <td>Default</td> <td>0 seconds</td> </tr> </table>	Values	1 — 600 seconds	Default	0 seconds
Values	1 — 600 seconds				
Default	0 seconds				

snmp

Syntax	snmp
Context	config>service>vprn
Description	This command enables 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 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	<p>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.</p> <p>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).</p> <p>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).</p> <p>The no form of the command removes the SNMP community name from the given VPRN context.</p>
Default	None — The SNMP community must be explicitly specified.
Parameters	<p><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).</p> <p>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.</p> <p>hash2 — Specifies the key is entered in a more complex encrypted form.</p> <p>version <i>SNMP-version</i> — Specifies the SNMP version.</p> <p>Values v1, v2c, both</p> <p><i>access-permissions</i> — Specifies the access rights to MIB objects.</p> <p>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.</p> <p><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.</p>

source-address

Syntax	source-address
Context	config>service>vprn
Description	This command enables 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> — Specify 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, etc.), the entire string must be enclosed within double quotes.

static-route

Syntax	[no] static-route { <i>ip-prefix/prefix-length</i> <i>ip-prefix netmask</i> } [preference <i>preference</i>] [metric <i>metric</i>] [tag <i>tag</i>] [community <i>comm-id</i>] [enable disable] { next-hop <i>ip-int-name</i> <i>ip-address</i> [mcast-family] ipsec-tunnel <i>ipsec-tunnel-name</i> } [bfd-enable { cpe-check <i>cpe-ip-address</i> [interval <i>seconds</i>] [drop-count <i>count</i>] [padding-size <i>padding-size</i>] [log]}] [no] static-route { <i>ip-prefix/prefix-length</i> <i>ip-prefix netmask</i> } [preference <i>preference</i>] [metric <i>metric</i>] [tag <i>tag</i>] [community <i>comm-id</i>] [enable disable] indirect <i>ip-address</i> [cpe- check <i>cpe-ip-address</i> [interval <i>seconds</i>][drop-count <i>count</i>] [padding-size <i>padding-size</i>] [log]] [no] static-route { <i>ip-prefix/prefix-length</i> <i>ip-prefix netmask</i> } [preference <i>preference</i>] [metric <i>metric</i>] [tag <i>tag</i>] [community <i>comm-id</i>] [enable disable] black-hole [mcast- family]
Context	config>service>vprn
Description	This command creates a static route. A static route can have a directly-connected interface as a next-hop (specified using the IP interface name or an IP address of the interface), or an indirect IP address as a next-hop or a black-hole next-hop (specifying a discard action). 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.

If a CPE connectivity check target address is already being used as the target address in a different static route, then `cpe-check` parameters must match. If they do not, the new configuration command will be rejected.

If a `static-route` command is issued with no `cpe-check` target but the destination prefix/netmask and next-hop matches a static route that did have an associated `cpe-check`, the `cpe-check` test will be removed from the associated static route.

Default No static routes are defined.

Parameters *ip-prefix* — The destination address of the aggregate route in dotted decimal notation.

Values	<code>ipv4-prefix</code>	a.b.c.d (host bits must be 0)
	<code>ipv4-prefix-length</code>	0 — 32

netmask — The subnet mask in dotted decimal notation.

Values	0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)
---------------	--

ip-int-name — The name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed with

ip-address — The IP address of the IP interface. The *ip-addr* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values	<code>ipv4-address</code>	a.b.c.d (host bits must be 0)
---------------	---------------------------	-------------------------------

community *comm-id* — 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.

Values	<code>comm-id</code>	<code>asn:comm-val well-known-comm</code>
	<code>asn</code>	0 — 65535
	<code>comm-val</code>	0 — 65535
	<code>well-known-comm</code>	<code>no-advertise, no-export, no-export-subconfed</code>

enable — Static routes can be administratively enabled or disabled. Use the **enable** parameter to re-enable a disabled static route. In order to enable a static route, it must be uniquely identified by the IP address, mask, and any other parameter that is required to identify the exact static route.

The administrative state is maintained in the configuration file.

Default	<code>enable</code>
----------------	---------------------

disable — Static routes can be administratively enabled or disabled. Use the **disable** parameter to disable a static route while maintaining the static route in the configuration. In order to enable a static route, it must be uniquely identified by the IP address, mask, and any other parameter that is required to identify the exact static route.

The administrative state is maintained in the configuration file.

Default	<code>enable</code>
----------------	---------------------

interval *seconds* — This optional parameter specifies the interval between ICMP pings to the target IP address.

Values 1 —255 seconds

Default 1 seconds

drop-count *count* — This optional parameter specifies the number of consecutive ping-replies that must be missed to declare the CPE down and to de-active the associated static route.

Values Value range: 1 —255

Default 3

padding-size *padding-size* — 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** parameter is used with the **static-route** command.

Values 0 — 16384 bytes

log — This optional parameter enables the ability to log transitions between active and in-active based on the CPE connectivity check. Events should be sent to the system log, syslog and SNMP traps.

next-hop [*ip-address* | *ip-int-name*] — Specifies the directly connected next hop IP address used to reach the destination. If the next hop is over an unnumbered interface, the *ip-int-name* of the unnumbered interface (on this node) can be configured.

The **next-hop** keyword and the **indirect** or **black-hole** keywords are mutually exclusive. If an identical command is entered (with the exception of either the **indirect** or **black-hole** parameters), then this static route will be replaced with the newly entered command, and unless specified, the respective defaults for preference and metric will be applied.

The *ip-addr* configured here can be either on the network side or the access side on this node. This address must be associated with a network directly connected to a network configured on this node.

ipsec-tunnel *ipsec-tunnel-name* — specifies an IPsec tunnel name up to 32 characters in length.

indirect *ip-address* — Specifies that the route is indirect and specifies the next hop IP address used to reach the destination.

The configured *ip-addr* is not directly connected to a network configured on this node. The destination can be reachable via multiple paths. The static route remains valid as long as the address configured as the indirect address remains a valid entry in the routing table. Indirect static routes cannot use an ip-prefix/mask to another indirect static route.

The **indirect** keyword and the **next-hop** or **black-hole** keywords are mutually exclusive. If an identical command is entered (with the exception of either the **next-hop** or **black-hole** parameters), then this static route will be replaced with the newly entered command and unless specified the respective defaults for preference and metric will be applied.

The *ip-addr* configured can be either on the network or the access side and is normally at least one hop away from this node.

black-hole — Specifies a black hole route meaning that if the destination address on a packet matches this static route it will be silently discarded.

The **black-hole** keyword is mutually exclusive with either the **next-hop** or **indirect** keywords. If an identical command is entered, with exception of either the **next-hop** or **indirect** parameters, then the static route is replaced with the new command, and unless specified, the respective defaults for **preference** and **metric** are applied.

preference *preference* — The preference of this static route (as opposed to the routes from different sources such as BGP or OSPF), expressed as a decimal integer. When modifying the **preference** value of an existing static route, unless specified, the metric will not change.

If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of which route to use is determined by the configuration of the ECMP command.

Default 5

Values 1 — 255

metric *metric* — The cost metric for the static route, expressed as a decimal integer. This value is used when importing this static route into other protocols such as OSPF. This value is also used to determine the static route to install in the forwarding table: When modifying the metrics of an existing static route, unless specified, the preference will not change.

If there are multiple static routes with the same preference but unequal metrics, the lower cost (metric) route is installed. If there are multiple static routes with equal preference and metrics then ECMP rules apply. If there are multiple routes with unequal preferences, then the lower preference route is installed.

Default 1

Values 0 — 65535

tag — Adds a 32-bit integer tag to the static route. The tag is used in route policies to control distribution of the route into other protocols.

Values 1 — 4294967295

bfd-enable — Associates the state of the static route to a BFD session between the local system and the configured nexthop. This keyword cannot be configured if the nexthop is **indirect** or a **blackhole** keywords are specified.

cpe-check *target-ip-address* — This parameter specifies the IP address of the target CPE device. ICMP pings will be sent to this target IP address. This parameter must be configured to enable the CPE connectivity feature for the associated static route. The target-ip-address cannot be in the same subnet as the static route subnet itself to avoid possible circular references. This option is mutually exclusive with BFD support on a given static route.

Default no cpe-check enabled

mcast-family — Enables submission of the IPv4 static route into IPv4 multicast RTM.

Values mcast-ipv4

ttl-propagate

Syntax	ttl-propagate
Context	config>service>vprn
Description	This command enables the context to configure TTL propagation for transit and locally generated packets in a given VPRN routing context.

Global Commands

Default none

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 inherit

Parameters **inherit** — The TTL propagation behavior is inherited from the global configuration under config>router>ttl-propagate>vprn-local
none — The TTL of the IP packet is not propagated into the VC label or labels in the transport label stack
vc-only — The TTL of the IP packet is propagated into the VC label and not into the labels in the transport label stack.
all — 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

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 config>router>ttl-propagate>vprn-transit.
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 inherit

Parameters **inherit** — The TTL propagation behavior is inherited from the global configuration under config>router>ttl-propagate>vprn-transit
none — The TTL of the IP packet is not propagated into the VC label or labels in the transport label stack
vc-only — The TTL of the IP packet is propagated into the VC label and not into the labels in the transport label stack.

all — 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	<p>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.</p> <p>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.</p> <p>subscriber-split-horizon — Controls the flow of traffic for wholesale subscriber applications.</p>

vrf-export

Syntax	vrf-export <i>policy</i> [<i>policy...</i>] no vrf-export
Context	config>service>vprn
Description	<p>This command specifies the export policies to control routes exported from the local VPN routing/forwarding (VRF) to other VRFs on the same or remote PE routers (via MP-BGP).</p> <p>You can specify up to fifteen (15) policy names.</p> <p>The no form of the command removes all route policy names from the export list.</p>
Default	None — No routes are exported from the VRF by default.
Parameters	<i>policy</i> — The route policy statement name.

vrf-import

Syntax	vrf-import <i>policy</i> [<i>policy...</i>] no vrf-import
Context	config>service>vprn

Description	<p>This command sets the import policies to control routes imported to the local VPN routing/forwarding (VRF) from other VRFs on the same or remote PE routers (via MP-BGP). Up to fifteen (15) names may be specified.</p> <p>BGP-VPN routes imported with a vrf-import policy will use the BGP preference value of 170 when imported from remote PE routers, or retain the protocol preference value of the exported route when imported from other VRFs on the same router, unless the preference is changed by the policy.</p> <p>The no form of the command removes all route policy names from the import list</p>
Default	None — No routes are accepted into the VRF by default.
Parameters	<i>policy</i> — The route policy statement name.

vrf-target

Syntax	vrf-target {<i>ext-community</i> export <i>ext-community</i> import <i>ext-community</i>} no vrf-target
Context	config>service>vpn
Description	<p>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).</p> <p>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.</p> <p>Specified vrf-import or vrf-export policies override the vrf-target policy.</p> <p>The no form of the command removes the vrf-target</p>
Default	no vrf-target
Parameters	<p><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.</p> <p>Values <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..65535] 2byte-asnumber [0..65535] ext-comm-val [0..4294967295] 4byte-asnumber [0..4294967295]</p> <p>import <i>ext-community</i> — Specify communities allowed to be accepted from remote PE neighbors.</p> <p>export <i>ext-community</i> — Specify communities allowed to be sent to remote PE neighbors.</p>

server

Syntax	server <i>server-name</i> no server
Context	configure>service>ies>sub-if>grp-if>local-address-assignment configure>service>ies>sub-if>local-address-assignment configure>service>vprn>sub-if>grp-if>local-address-assignment configure>service>vprn>sub-if>local-address-assignment
Description	This command designates a local 7x50 DHCPv4 server for local pools management where IPv4 addresses for PPPoXv4 clients will be allocated without the need for the internal 7x50 DHCP relay-agent. Those addresses will be tied to PPPoX sessions and they will be de-allocated when the PPPoX session is terminated.
Default	none
Parameters	<i>server-name</i> — Specifies the name of the local 7x50 DHCP server.

client-application

Syntax	client-application [ppp-v4] no client-application
Context	configure>service>ies>sub-if>grp-if>local-address-assignment configure>service>ies>sub-if>local-address-assignment configure>service>vprn>sub-if>grp-if>local-address-assignment configure>service>vprn>sub-if>local-address-assignment
Description	This command enables local 7x50 DHCP Server pool management for PPPoXv4 clients. A pool of IP addresses can be shared between IPoE 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	none

default-pool

Syntax	default-pool <i>pool-name</i> no default-pool configure>service>ies>sub-if>grp-if>local-address-assignment configure>service>ies>sub-if>local-address-assignment configure>service>vprn>sub-if>grp-if>local-address-assignment configure>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.
Default	none

Parameters *pool-name* — Specifies the name of the local 7x50 DHCP server pool.

delayed-enable

Syntax	delayed-enable <i>seconds</i> [init-only] no delayed-enable
Context	configure>service>ies>sub-if>local-address-assignment configure>service>vprn>sub-if>local-address-assignment
Description	<p>This command will render the subscriber-interface non operation for the given amount of time once the node is rebooted or once the interface is enabled (no-shutdown). The purpose of this timer is to stall the operation of the subscriber-interface until the MCS database is synchronized.</p> <p>A typical use case for this timer would be to prevent IP lease duplication for PPPoE clients using local PPPoXv4/v6 pools in redundant DHCPv4/v6 server configuration. Since there is no classical DHCP lease state maintained for local PPPoXv4/v6 pools, the IP addresses will not be synchronized via DHCP Server. Instead they will be synchronized via PPPoX clients whose state is maintained in 7x50. Once the PPPoX subscriber host is synchronized between the two 7x50 nodes, the respective IP address lease will be updated in the respective local pool.</p> <p>One artifact of this behavior (IP address assignment in local DHCP pools is synchronized via PPPoX clients and not via DHCP server synchronization mechanism) is that during the node boot, the DHCP server must wait for the completion of PPPoX subscriber synchronization via MCS so that it learns which addresses/prefixes are already allocated on the peering node. Since the DHCP server can theoretically start assigning IP addresses before the PPPoX sync is completed, a duplicate address assignment may occur. For example an IP address lease can be granted via DHCP local pools while PPPoX sync is still in progress. Once the PPPoX sync is completed, the DHCP server may discover that the granted IP lease has already been allocated by the peering node. The most recent lease will be kept and the other will be removed from both systems. To prevent this scenario, a configurable timer is set to an arbitrary value that will render sub-if non-operational until the timer expires. The purpose of this timer is to allow the PPPoX sync to complete before subscribers under the sub-intf can be served.</p>
Default	none
Parameters	<i>second</i> — Specifies in seconds.
Values	1 — 1200

Log Commands

log

Syntax	log
Context	configure>service>vprn configure>service>vprn>log-id
Description	This command enables 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	<p>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.</p> <p>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.</p> <p>Any changes made to an existing filter, using any of the sub-commands, are immediately applied to the destinations where the filter is applied.</p> <p>The no form of the command removes the filter association from log IDs which causes those logs to forward all events.</p>
Default	No event filters are defined.
Parameters	<i>filter-id</i> — The filter ID uniquely identifies the filter.
Values	1 — 1000

default-action

Syntax	default-action {drop forward} no default-action
Context	config>service>vprn>log>filter
Description	<p>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.</p> <p>When multiple default-action commands are entered, the last command overwrites the previous command.</p>

Log Commands

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 — 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 **application** 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 context is defined.

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 etc. 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. Valid operators are listed in the table below.

Operator	Notes
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, vtr

message

- Syntax** **message** {**eq** | **neq**} **pattern** *pattern* [**regex**]
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.
regex — Specifies the type of string comparison to use to determine if the log event matches the value of **message** command parameters. When the **regex** 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 in the table below. Valid operators are:

Operator	Notes
eq	equal to
neq	not equal to
lt	less than

Operator	Notes
lte	less than or equal to
gt	greater than
gte	greater than or equal to

event-id — The event ID, expressed as a decimal integer.

Values 1 — 4294967295

severity

Syntax	severity { eq neq lt lte gt gte } <i>severity-level</i> 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 — No severity level match criterion is specified.
Parameters	eq neq lt lte gt gte — This operator specifies the type of match. Valid operators are listed in the table below.

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. The following table lists severity names and corresponding numbers per ITU standards M.3100 X.733 & X.21 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 [regexp]**
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** — No subject match criterion specified.

Parameters **eq | neq** — This operator specifies the type of match. Valid operators are listed in the following table:

Operator	Notes
eq	equal to
neg	not equal to

subject — A string used as the subject match criterion.

regexp — Specifies the type of string comparison to use to determine if the log event matches the value of **subject** command parameters. When the **regexp** 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 **regexp** 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	<p>This command creates a context to configure destinations for event streams.</p> <p>The log-id context is used to direct events, alarms/traps, and debug information to respective destinations.</p> <p>A maximum of 10 logs can be configured.</p> <p>Before an event can be associated with this log-id, the from command identifying the source of the event must be configured.</p> <p>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.</p> <p>Use the event-control command to suppress the generation of events, alarms, and traps for all log destinations.</p> <p>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.</p> <p>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.</p> <p>Note that 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.</p> <p>The no form of the command deletes the log destination ID from the configuration.</p>
Default	No log destinations are defined.
Parameters	<i>log-id</i> — The log ID number, expressed as a decimal integer.
Values	1 — 100

to snmp

Syntax	to snmp [<i>size</i>]
Context	config>service>vprn>log>log-id
Description	<p>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>.</p> <p>A local circular memory log is always maintained for SNMP notifications sent to the specified snmp-trap-group for the <i>log-id</i>.</p> <p>The source of the data stream must be specified in the from command prior to configuring the destination with the to command.</p> <p>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.</p>

Log Commands

Default none

Parameters *size* — The *size* parameter defines the number of events stored in this memory log.

Default 100

Values 50 — 1024

to syslog

Syntax **to syslog** *syslog-id*

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.

Default none

Parameters *syslog-id* — Instructs the events selected for the log ID to be directed to the *syslog-id*. The characteristics of the *syslog-id* referenced here must have been defined in the **config>log>syslog** *syslog-id* context.

Values 1 — 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 *log-id*. 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 *log-id*. 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 *log-id*. 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** 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** *syslog-id*
- 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 *syslog-id* 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** *syslog-id* — The syslog ID number for the syslog destination, expressed as a decimal integer.
- Values** 1 — 10

address

Syntax	address <i>ip-address</i> no address
Context	config>service>vprn>log>syslog
Description	<p>This command adds the syslog target host IP address to/from a syslog ID.</p> <p>This parameter is mandatory. If no address is configured, syslog data cannot be forwarded to the syslog target host.</p> <p>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.</p> <p>The same syslog target host can be used by multiple log IDs.</p> <p>The no form of the command removes the syslog target host IP address.</p>
Default	no address — There is no syslog target host IP address defined for the syslog ID.
Parameters	<i>ip-address</i> — The IP address of the syslog target host in dotted decimal notation.

facility

Syntax	facility <i>syslog-facility</i> no facility
Context	cconfig>service>vprn>log>syslog
Description	<p>This command configures the facility code for messages sent to the syslog target host.</p> <p>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 <i>facility-code</i> entered overwrites the previous facility-code.</p> <p>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.</p> <p>The no form of the command reverts to the default value.</p>
Default	local7 — syslog entries are sent with the local7 facility code.
Parameters	<p><i>syslog-facility</i> — 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.</p> <p>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</p>

Valid responses per RFC3164, *The BSD syslog Protocol*, are listed in the table below.

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 — 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 a 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.

Log Commands

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 **no log-prefix** — no prepend log prefix string defined.

Parameters *log-prefix-string* — An alphanumeric string of up to 32 characters. Spaces and colons (:) cannot be used in the string.

level

Syntax **level** *syslog-level*
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.

Parameters *value* — 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
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	<p>This command configures the UDP port that will be used to send syslog messages to the syslog target host.</p> <p>The port configuration is needed if the syslog target host uses a port other than the standard UDP syslog port 514.</p> <p>Only one port can be configured. If multiple port commands are entered, the last entered port overwrites the previously entered ports.</p> <p>The no form of the command reverts to default value.</p>
Default	no port
Parameters	<p><i>value</i> — The value is the configured UDP port number used when sending syslog messages.</p> <p>Values 1 — 65535</p>

snmp-trap-group

Syntax	[no] snmp-trap-group <i>log-id</i>
Context	config>service>vprn>log
Description	<p>This command creates the context to configure a group of SNMP trap receivers and their operational parameters for a given log-id.</p> <p>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.</p> <p>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.</p> <p>The no form of the command deletes the SNMP trap group.</p>
Default	There are no default SNMP trap groups.
Parameters	<p><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.</p> <p>Values 1 — 99</p>

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 }] no trap-target <i>name</i>
Context	config>service>vprn>log>snmp-trap-group
Description	<p>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.</p> <p>Before an SNMP trap can be issued to a trap receiver, the log-id, snmp-trap-group and at least one trap-target must be configured.</p> <p>The trap-target command is used to add/remove a trap receiver from an snmp-trap-group. The operational parameters specified in the command include:</p> <ul style="list-style-type: none"> • 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. <p>A single snmp-trap-group <i>log-id</i> can have multiple trap-receivers. Each trap receiver can have different operational parameters.</p> <p>An address can be configured as a trap receiver more than once as long as a different port is used for each instance.</p> <p>To prevent resource limitations, only configure a maximum of 10 trap receivers.</p> <p>Note that if the same trap-target <i>name</i> port <i>port</i> parameter value is specified in more than one SNMP trap group, each trap destination should be configured with a different <i>notify-community</i> 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.</p> <p>The no form of the command removes the SNMP trap receiver from the SNMP trap group.</p>
Default	No SNMP trap targets are defined.
Parameters	<p><i>name</i> — Specifies the name of the trap target up to 28 characters in length.</p> <p>address <i>ip-address</i> — The IP address of the trap receiver in dotted decimal notation. Only one IP address destination can be specified per trap destination group.</p> <p>Values <i>ipv4-address</i> a.b.c.d (host bits must be 0)</p> <p>port <i>port</i> — 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.</p> <p>Default 162</p> <p>Values 1 — 65535</p>

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.

Default `snmpv3`

Values `snmpv1, snmpv2c, 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**.

Default no-auth-no-privacy. This parameter can only be configured if SNMPv3 is also configured.

Values no-auth-no-privacy, auth-no-privacy, privacy

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. Note that 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.

Multicast VPN Commands

mvpn

Syntax	mvpn
Context	config>service>vprn
Description	This command enables the context to configure MVPN-related parameters for the IP VPN.

auto-discovery

Syntax	[default mdt-safi] [source-address <i>ip-address</i>]
Context	config>service>vprn>mvpn
Description	<p>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.</p> <p>The no form of the command disables MVPN membership auto-discovery through BGP.</p>
Default	<p>default</p> <p>default — Enable AD route exchange based on format defined in ng-MVPN (RFC6514).</p> <p>mdt-safi — Enable AD route exchange based on mdt-safi format defined in <i>draft-rosen-vpn-mcast</i>.</p> <p>This command allows optionally to specify a source-address - an IP address to be used by Rosen M-VPN for core diversity non-default IGP instances (not using system IP). Two unique IP addresses for all MVPNs are supported. For instances using default System IP, source address configuration should not be specified to avoid consuming one of the addresses.</p> <p>Explicitly defined source-address allows GRE-encapsulated Rosen MVPN multicast traffic (Default and Data MDT) to originate from a configured IP address, so the source IP address of the GRE packets won't be the default system IP address.</p> <p>Value:</p> <p><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.</p>

c-mcast-signaling

Syntax	c-mcast-signaling {bgp pim} no c-mcast-signaling
Context	config>service>vprn>mvpn

Log Commands

Description	<p>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.</p> <p>Changes may only be made to this command when the mvpn node is shutdown.</p> <p>The no form of the command reverts it back to the default.</p>
Default	mcast-signaling bgp
Parameters	<p>bgp — Specifies to use BGP for PE-to-PE signaling of CEmulticast states. Auto-discovery must be enabled.</p> <p>pim — Specifies to use PIM for PE-to-PE signaling of CE multicast states.</p>

intersite-shared

Syntax	intersite-shared [persistend-type5-adv] [kat-type5-adv-withdraw] no intersite-shared
Context	config>service>vprn>mvpn
Description	<p>This command specifies whether to use inter-site shared C-trees or not. Optional parameters allow enabling additional intersite-shared functionality as per below description. Not specifying an optional parameter when executing the command disables that parameter.</p>
Default	intersite-shared
Parameters	<p>persistend-type5-adv — when specified for intersite-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.</p> <p>kat-type5-adv-withdraw — when specified for intersite-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 withdraw 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. Note the following:</p> <ul style="list-style-type: none">• 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	<p>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.</p> <p>This command provides an option to configure either a sender-only or receiver only mode per PE node. Restricting the role of a PE node avoids creating full mesh of MDT/PMSI across all PE nodes that are participating in MVPN instance</p> <p>The no version of this command restores the default (sender-receiver).</p>
Default	mdt-type sender-receiver
Parameters	<p>sender-receiver — MVPN has both sender and receivers connected to PE node</p> <p>sender-only — MVPN has only senders connected to PE node</p> <p>receiver-only — MVPN has only receivers connected to PE node</p>

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	<p>src-prefix <i>ip-address/mask</i> [<i>ip-address/mask ...up to 8 maximum</i>]</p> <p>no src-prefix <i>ip-address/mask</i></p>
Context	config>service>vprn>mvpn>red-source-list
Description	<p>This command configures up to 8 multicast source IPv4 prefixes for preferred source selection. Single or multi-line inputs are allowed.</p> <p>The no form of the command deletes specified prefix from the list.</p>
Default	no prefixes are specified
Parameters	<i>ip-address/mask</i> — IPv4 address prefix with mask.

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 ...up to 8 maximum</i>] no <i>ipv6-ip-address/prefix-length</i>
Context	config>service>vprn>mvpn>red-source-list>ipv6
Description	This command configures up to 8 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.

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...(upto 8 max)</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.

Default none

Parameters *ip-address/mask* — Specifies the IPv4 multicast address prefix with mask.

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>pt

Description This command enables the context for specifying inclusive provider tunnels

bsr

Syntax **bsr {unicast | spmsi}**
no bsr

Context config>service>vprn>mvpn>provider-tunnel>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 PDU's are sent/forwarded using unicast PDUs (default).
spmsi — BSR PDU's are sent/forwarded using S-PMSI full mesh.

mldp

Syntax **mldp**
no mldp

Context config>service>vprn>mvpn>provider-tunnel>inclusive

Description This command enables use of mLDP LSP for the provider tunnel.

Log Commands

Default no mldp

shutdown

Syntax **shutdown**
no shutdown

Context config>service>vprn>mvpn>provider-tunnel>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.

Note that 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 use PIM SSM for inclusive provider tunnels.

group-address — Specifies the multicast group address to use.

hello-interval

Syntax **hello-interval hello-interval**
no hello-interval

Context config>service>vprn>mvpn>provider-tunnel>inclusive>pim

Description This command specifies the interval at which PIM hello messages are transmitted on the PIM inclusive provider tunnel.

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

Default 30 seconds

Parameters *hello-interval* — Specifies the hello interval, in seconds. A 0 (zero) value disables the sending of hello messages.

Values 0 — 255

hello-multiplier

Syntax	hello-multiplier <i>deci-units</i> no hello-multiplier
Context	config>service>vprn>mvpn>provider-tunnel>inclusive>pim
Description	This command specifies the hello multiplier. The hello-multiplier in conjunction with the hello-interval determines the hold time for a PIM neighbor. Hold time = (hello-interval * hello-multiplier) / 10. The no form of the command reverts the value to the default.
Default	35
Parameters	<i>deci-units</i> — Specifies the value, in multiples of 0.1, for the formula used to calculate the hold time
	Values 20 — 100

improved-assert

Syntax	[no] improved-assert
Context	config>service>vprn>mvpn>provider-tunnel>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>provider-tunnel>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>provider-tunnel>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

wildcard-spmsi

Syntax	wildcard-spmsi no wildcard-spmsi
Context	config>service>vprn>mvpn>provider-tunnel>inclusive
Description	This command enables RFC6625 (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 I-PMSI LSP template. 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 enables the context to specify selective provider tunnel parameters.
Default	none

auto-discovery-disable

Syntax	[no] auto-discovery-disable
Context	config>service>vprn>mvpn>provider-tunnel>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.

The **no** form of the command enables multicast VPN membership auto-discovery through BGP.

Default no auto-discovery-disable

data-delay-interval

Syntax	data-delay-interval <i>value</i> no data-delay-interval
Context	config>service>vprn>mvpn>provider-tunnel>selective
Description	This command specifies the interval, in seconds, before a PE router connected to the source switches traffic from the inclusive provider tunnel to the selective provider tunnel. The no form of the command reverts the value to the default.
Default	3 seconds
Parameters	<i>value</i> — Specifies the data delay interval, in seconds.
Values	3 — 180

data-threshold

Syntax	data-threshold { <i>c-grp-ip-addr/mask</i> <i>c-grp-ip-addr netmask</i> } <i>s-pmsi-threshold</i> [pe-threshold-add <i>pe-threshold-add</i>] [pe-threshold-delete <i>pe-threshold-delete</i>] data-threshold <i>c-grp-ipv6-addr/prefix-length</i> <i>s-pmsi-threshold</i> [pe-threshold-add <i>pe-threshold-add</i>] [pe-threshold-delete <i>pe-threshold-delete</i>] no data-threshold { <i>c-grp-ip-addr/mask</i> <i>c-grp-ip-addr netmask</i> } no data-threshold <i>c-grp-ipv6-addr/prefix-length</i>
Context	config>service>vprn>mvpn>provider-tunnel>selective
Description	This command specifies the data rate threshold that triggers the switch from the inclusive provider tunnel to the selective provider tunnel for (C-S, C-G) within the group range. Optionally, PE thresholds for creating/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. The no form of the command removes the values from the configuration.
Default	no data-threshold
Parameters	<i>group-address/mask</i> — Specifies a multicast group address and netmask length. <i>c-grp-ip-addr/mask</i> <i>c-grp-ip-addr netmask</i> — Specifies an IPv4 multicast group address and netmask length or network mask. <i>c-grp-ipv6-addr/prefix-length</i> — Specifies an IPv6 multicast group address and prefix length.

- s-pmsi-threshold* — Specifies the rate, in kilobits per second (kbps). 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..32]
- netmask* : a.b.c.d (network bits all 1 and host bits all 0)
- s-pmsi-threshold* : [1..4294967294](threshold in kbps)
- 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..FFFF]H
d [0..255]D
prefix-length [1..128]
- pe-threshold-add: [1..65535], if never specified, 65535 is used
(add threshold always met)
- pe-threshold-delete: [2..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>provider-tunnel>selective>pim

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

pim-asm

Syntax [no] pim-asm {grp-ip-address/mask | grp-ip-address netmask}

Context config>service>vprn>mvpn>provider-tunnel>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.

— enable-asm-mdt

Syntax	[no] enable-asm-mdt
Context	config>service>vprn>mvpn>provider-tunnel>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>provider-tunnel>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 ip-address standby ip-address no umh-pe ip-address
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 IP address is selected as UMH. hash-based — Specifies that the UMH selection is based on the hash based procedures. tunnel-status — Specifies that UMH selection is based on the state of the tunnel as well as the available unicast routes through the tunnel. unicast-rt-pref — When selected, best unicast route will decide which UMH is chosen. Note that, 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>policy-name</i> [<i>policy-name...</i> (up to 16 max)]} no vrf-export
Context	config>service>vprn>mvpn
Description	This command specifies the export policy (up to 16) 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>policy</i> — Specifies a route policy name.

vrf-import

Syntax	vrf-import { unicast <i>policy-name</i> [<i>policy-name...</i> (up to 16 max)]} no vrf-import
Context	config>service>vprn>mvpn
Description	This command specifies the import policy (up to 16) to control MVPN routes imported to the local VRF from other VRFs on the same or remote PE routers.
Default	vrf-import unicast

Parameters **unicast** — Specifies to use a unicast VRF import policy for the MVPN.
policy — Specifies a route policy name.

vrf-target

Syntax **vrf-target** {**unicast** | *ext-community* | **export unicast** | *ext-community* | **import unicast** | *ext-community*}
no vrf-target

Context config>service>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.
ext-comm — An extended BGP community in the **type:x:y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. **x** and **y** are 16-bit integers.

Values target: {*ip-address:comm-val* | *2byte-asnumber:ext-comm-val* | *4byte-asnumber:comm-val*}
ip-address: a.b.c.d
comm-val: 0 — 65535
2byte-asnumber: 1 — 65535
4byte-asnumber 0 — 4294967295

import *ext-community* — Specify communities allowed to be accepted from remote PE neighbors.
export *ext-community* — Specify communities allowed to be sent to remote PE neighbors.

export

Syntax **export** {**unicast** | *ext-community*}

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.
ext-comm — An extended BGP community in the **type:x:y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. **x** and **y** are 16-bit integers.

Values target: {*ip-address:comm-val* | *2byte-asnumber:ext-comm-val* | *4byte-asnumber:comm-val*}
ip-address: a.b.c.d
comm-val: 0 — 65535

<i>2byte-asnumber:</i>	1 — 65535
<i>4byte-asnumber</i>	0 — 4294967295

import

Syntax `import {unicast | ext-community}`

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.

ext-comm — An extended BGP community in the **type:x:y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. **x** and **y** are 16-bit integers.

Values target: {*ip-address:comm-val* | *2byte-asnumber:ext-comm-val* | *4byte-asnumber:comm-val*}

<i>ip-address:</i>	a.b.c.d
<i>comm-val:</i>	0 — 65535
<i>2byte-asnumber:</i>	1 — 65535
<i>4byte-asnumber</i>	0 — 4294967295

Note: des is not permitted in FIPS-140-2 mode.

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, etc.), the entire string must be enclosed within 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.

SDP Commands

spoke-sdp

Syntax	[no] spoke-sdp <i>sdp-id</i>
Context	config>service>vprn
Description	<p>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.</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>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 <i>sdp-id</i> is not already configured, an error message is generated. If the <i>sdp-id</i> does exist, a binding between that <i>sdp-id</i> and the service is created.</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 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.</p>
Default	No <i>sdp-id</i> is bound to a service.
Special Cases	VPRN — Several SDPs can be bound to a VPRN service. Each SDP must be destined to a different router. If two <i>sdp-id</i> bindings terminate on the same router, an error occurs and the second SDP binding is rejected.
Parameters	<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>Values 1 — 4294967295</p> <p>Values</p>

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 enables 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: <ul style="list-style-type: none"> • The system is using network chassis mode D • SDP signaling is off • control-word is enabled (control-word is disabled by default) • the service type is epipe, vpls, 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	<i>agi</i> — Specifies the attachment group identifier. Values 0 — 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	<p><i>global-id</i> — Specifies the global ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.</p> <p>Values 0 — 4294967295</p> <p><i>node-id</i> — Specifies the node ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.</p> <p>Values a.b.c.d or 0 — 4294967295</p> <p><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.</p> <p>Values 1 — 4294967295</p>

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	<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>Values 0 — 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>Values a.b.c.d or 0 — 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>Values 1 — 4294967295</p>

spoke-sdp

Syntax	spoke-sdp <i>sdp-id</i> [: <i>vc-id</i>] <i>vc-type</i> { ether ipipe } [create] no spoke-sdp <i>sdp-id</i> [: <i>vc-id</i>] <i>vc-type</i> { ether ipipe } [create]
Context	config>service>vprn>if
Description	<p>This command binds a service to an existing Service Distribution Point (SDP).</p> <p>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>The SDP must already be defined in the config>service>sdp context in order to associate an SDP with a service. If the sdp <i>sdp-id</i> is not already configured, an error message is generated. If the <i>sdp-id</i> does exist, a binding between that <i>sdp-id</i> and the service is created.</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>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.</p> <p>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.</p>
Default	none
Special Cases	<p>VPRN — Several SDPs can be bound to a VPRN service. Each SDP must be destined to a different router. If two <i>sdp-id</i> bindings terminate on the same , an error occurs and the second SDP is binding is rejected.</p> <p><i>sdp-id</i> — The SDP identifier.</p> <p>Values 1 — 17407</p> <p><i>vc-id</i> — The virtual circuit identifier.</p> <p>Values 1 — 4294967295</p> <p><i>vc-type</i> — The encapsulation and pseudowire type for the spoke SDP.</p> <p>Values ether—Ethernet pseudowire. ipipe—IP pseudowire.</p> <p>Default ether</p>

egress

Syntax	egress
Context	config>service>vprn>if>spoke-sdp

SDP Commands

```
config>service>vprn>red-if>spoke-sdp
```

Description This command configures an SDP context.

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 LDP or RSVP SDP as well as to a VPRN service using the autobind mode with the with the ldp, rsvp-te, or mpls options. 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. Note however that 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 <i>network-policy-id</i> fp-redirect-group <i>queue-group-name</i> instance <i>instance-id</i> no qos
Context	config>service>vprn>interface>spoke-sdp>ingress
Description	<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"> 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). 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. 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. 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. 5. One or more spoke-SDPs can have their FCs redirected to use policers in the same policer queue-group instance. <p>The following are the constraints and rules of this provisioning model when used in the ingress pseudowire rate-limiting feature:</p> <ol style="list-style-type: none"> 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. 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-SPD 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. 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.
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.
 5. If no network QoS policy is applied to the ingress context of the pseudowire, then 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/VP RN spoke interface and from an R-VPLS spoke-SPD, 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 — 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 — 16384

vc-label

Syntax	vc-label <i>egress-vc-label</i> no vc-label [<i>egress-vc-label</i>]
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	<i>vc-label</i> — A VC egress value that indicates a specific connection. Values 16 — 1048575

vc-label

Syntax	vc-label <i>ingress-vc-label</i> no vc-label [<i>ingress-vc-label</i>]
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	<i>vc-label</i> — A VC ingress value that indicates a specific connection. Values 2048 — 18431

egress

Syntax	egress
Context	config>service>vprn>network-interface
Description	This command enables the context to configure egress network filter policies for the interface.

filter

Syntax	filter ip <i>ip-filter-id</i> no filter [<i>ip ip-filter-id</i>] [ipv6 <i>ipv6-filter-id</i>]
Context	config>service>vprn>network-interface>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

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

Values 1 — 65535

qos

Syntax **qos** *network-policy-id* **port-redirect-group** *queue-group-name* [**instance** *instance-id*]
no qos [*network-policy-id*]

Context configure>service>apipe>spoke-sdp>egress
configure>service>cpipe>spoke-sdp>egress
configure>service>epipe>spoke-sdp>egress
configure>service>fpipe>spoke-sdp>egress
configure>service>ipipe>spoke-sdp>egress
config>service>vpls>spoke-sdp>egress
config>service>vpls>mesh-sdp>egress
config>service>pw-template>egress
config>service>vprn>interface>spoke-sdp>egress
config>service>ies>interface>spoke-sdp>egress

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-SPD inside a service or to the egress context of a pseudowire template and specify the redirect queue-group name.

One or more spoke-SPDs 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-SPD 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-SPD 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:
 - 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.
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/dot1.p/DSCP and the tunnel DEI/dot1.p/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/dot1.p and the tunnel DEI/dot1.p/EXP, but the DSCP is not modified by the policer operation.

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/dot1.p/DSCP and the tunnel DEI/dot1.p/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 — 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 — 16384

Interface Commands

interface

Syntax	interface <i>ip-int-name</i> no interface <i>ip-int-name</i>
Context	config>service>vprn
Description	<p>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.</p> <p>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 internet access.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p>
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, etc.), the entire string must be enclosed within double quotes.

Interface Commands

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.

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 { <i>ip-address/mask</i> <i>ip-address netmask</i> } [broadcast <i>all-ones</i> host-ones] [track-srrp <i>srrp-instance</i>] no address
Context	config>service>vprn>if config>service>vprn>nw-if
Description	<p>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.</p> <p>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 .</p> <p>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.</p> <p>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.</p> <p>By default, no IP address or subnet association exists on an IP interface until it is explicitly created.</p> <p>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.</p>

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.

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. Note that 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. Note that 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	<p>This command controls the forwarding of directed broadcasts out of the IP interface.</p> <p>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.</p> <p>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.</p> <p>When disabled, directed broadcast packets discarded at this egress IP interface will be counted in the normal discard counters for the egress SAP.</p> <p>By default, directed broadcasts are not allowed and will be discarded at this egress IP interface.</p> <p>The no form of this command disables the forwarding of directed broadcasts out of the IP interface.</p>
Default	no allow-directed-broadcasts — Directed broadcasts are dropped.

bfd

Syntax	bfd <i>transmit-interval</i> [receive <i>receive-interval</i>] [multiplier <i>multiplier</i>] [echo-receive <i>echo-interval</i>] no bfd
Context	config>service>vprn>if config>service>vprn>if>ipv6 config>service>vprn>nw-if

Description	<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 no form of the command removes BFD from the associated IGP protocol adjacency.</p> <p>Important notes: On the and 7950 XRS, the <i>transmit-interval</i>, receive <i>receive-interval</i>, and echo-receive <i>echo-interval</i> values can only be modified to a value less than 100 when:</p> <ol style="list-style-type: none"> 1. The type cpm-np option is explicitly configured. 2. The service is shut down (shutdown) 3. The interval is specified 10 — 100000. 4. The service is re-enabled (no shutdown) <p>To remove the type cpm-np option, re-issue the bfd command without specifying the type parameter.</p>
Default	no bfd
Parameters	<p><i>transmit-interval</i> — Sets the transmit interval for the BFD session.</p> <p>Values 10 — 100000</p> <p>Default 100</p> <p><i>receive</i> <i>receive-interval</i> — Sets the receive interval for the BFD session.</p> <p>Values 10 — 100000</p> <p>Default 100</p> <p>multiplier <i>multiplier</i> — Set the multiplier for the BFD session.</p> <p>Values 3— 20</p> <p>Default 3</p> <p>echo-receive <i>echo-interval</i> — Sets the minimum echo receive interval, in milliseconds, for the BFD session.</p> <p>Values 100 — 100000</p> <p>Default 100</p>

cflowd-parameters

Syntax	cflowd-parameters no cflowd-parameters
Context	config>service>vprn>if config>service>vprn>nw-if
Description	This command creates the configuration context to configure cflowd parameters for the associated IP interfaces.

Interface Commands

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

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 —

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 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.

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	not enabled

dist-cpu-protection

Syntax	dist-cpu-protection <i>policy-name</i> no dist-cpu-protection
Context	config>service>vprn>if>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 SAP or a network interface (note that this rule does not apply to templates such as an msap-policy)
Default	no dist-cpu-protection

delayed-enable

Syntax	delayed-enable <i>seconds</i> no delayed-enable
Context	config>service>vprn>if config>service>vprn>if>nw-if config>service>vprn>if>redundant-if
Description	This command will cause a delay in the activation of an IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface up. 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 expires.
Parameters	<i>seconds</i> — Specifies a delay, in seconds, to make the interface operational.
Values	1 — 1200

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 <template-name> 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> — Name of template (32 characters maximum)

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.
Default	none

dns

Syntax	dns ip-address [secondary ip-address] dns secondary ip-address no dns [<i>ip-address</i>] [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> — This parameter 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 — This parameter 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

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> [preferred] no link-local-address
Context	config>router>if>ipv6 config>service>ies>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. Note that 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	preferred — Disables duplicated address detection and sets the address to preferred, even if there is a duplicate 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.

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	<i>ipv6-address</i> x:x:x:x:x:x:x (eight 16-bit pieces) 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 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.

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 — 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 — 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 {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 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. The no form of this command includes both source and destination parameters.
Default	no egr-ip-load-balancing
Parameters	source — Specifies using the source address and (if l4-load balancing is enabled) source port in the hash, ignore destination address/port. destination — Specifies using the destination address and (if l4-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.

lsr-load-balancing

Syntax	lsr-load-balancing hashing-algorithm 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.

Interface Commands

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	disabled

teid-load-balancing

Syntax	[no] teid-load-balancing
Context	config>service>vprn>interface>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	disabled

local-proxy-arp

Syntax	[no] local-proxy-arp
Context	config>service>vprn>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 None

mac

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

Context config>service>vprn>if
config>service>vprn>if>vrrp
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.

ntp-broadcast

Syntax [no] ntp-broadcast

Context config>service>vprn>nw-if

Description This command enables receiving of NTP/SNTP broadcasts on the interface/

monitor-oper-group

Syntax	monitor-oper-group <i>name</i> 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	<i>name</i> — 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> (up to 5 max)]
Context	config>service>vprn>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
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, etc.), the entire string must be enclosed within double quotes. The specified name(s) must already be defined.

ptp-hw-assist

Syntax	[no] ptp-hw-assist
Context	config>service>vprn>if

Description	<p>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.</p> <p>If the SAP configuration of the interface is removed, the ptp-hw-assist configuration will be removed.</p> <p>If the IPv4 address configuration of the interface is removed, the ptp-hw-assist configuration will be removed.</p> <p>Only one interface per physical port can have ptp-hw-assist enabled.</p>
Default	no ptp-hw-assist

qos-route-lookup

Syntax	<p>qos-route-lookup [source destination] no qos-route-lookup</p>
Context	<pre>config>service>vprn>if config>service>vprn>if>ipv6</pre>
Description	<p>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.</p> <p>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.</p> <p>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.</p> <p>If neither the optional source or destination parameter is present, then the default is destination address matching.</p> <p>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.</p> <p>The no form of the command reverts to the default.</p>
Default	destination
Parameters	<p>source — Enables QoS classification of incoming IP packets based on the source address matching a route with QoS information.</p> <p>destination — Enables QoS classification of incoming IP packets based on the destination address matching a route with QoS information.</p>

redundant-interface

Syntax	redundant-interface <i>red-ip-int-name</i> no redundant-interface
Context	config>service>vprn
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>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 an secondary IP address/IP subnet/broadcast address format to the interface.
Default	none
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). <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. Note that 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.

static-arp

Syntax	static-arp <i>ieee-mac-address unnumbered</i> no static-arp unnumbered
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.
Default	none
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	<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>
Default	none
Parameters	<i>ip-address</i> — Specifies the static ISA tunnel redundant next-hop address.

strip-label

Syntax	[no] strip-label
Context	config>service>vprn>nw-if
Description	<p>This command forces packets to be stripped of all (max 5) MPLS labels before the packets are handed over for possible filter (PBR) processing.</p> <p>If the packets do not have an IP header ediatly following the MPLS label stack after the strip, they are discarded. Only MPLS encapsulated IP, IGP shortcuts and VPRN over MPLS packets will be processed.</p> <p>This command is only supported on:</p> <ul style="list-style-type: none">• Optical ports• IOM3-XP cards• Null/Dot1q encaps• Network ports• IPv4 <p>The no form removes the strip-label command.</p> <p>In order to associate an interface that is configured with the strip-label parameter with a port, the port must be configured as single-fiber for the command to be valid.</p>

Default no strip-label

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 *modifier*
[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 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>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).

Interface Commands

Parameters *bits* — Specifies the number of bits.

Values 512–1024

security-parameter

Syntax **security-parameter** *sec*
[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 *sec* — Specifies the security parameter.

Values 0–1

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 - 65535

tcp-mss

Syntax	tcp-mss <i>mss-value</i> 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	<i>mss-value</i> — 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 = max-IP_MTU (9198)-40 Values 536 - 9158 (IPv4) 1220 - 9138 (IPv6)

tos-marking-state

Syntax	tos-marking-state {trusted untrusted} no tos-marking-state
Context	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.

Interface Commands

Default 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.

unnumbered

Syntax **unnumbered** [*ip-int-name* | *ip-address*]
no unnumbered

Context config>service>vprn>if
config>service>vprn>nw-if

Description This command configures the interface as an unnumbered interface. Unnumbered IP interface is supported on a Sonet/SDH access port with the PPP, ATM, or Frame Relay encapsulation. It is not supported on a TDM port or channel.

Parameters *ip-int-name* — Specifies the name of an IP interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

ip-address — Specifies an IP address.

qos

Syntax **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

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 — 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 — 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 queue-group-name 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 — 16384

urpf-check

Syntax	[no] urpf-check
Context	config>service>vprn>if config>service>vprn>nw-if config>service>vprn>interface>ipv6
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	disabled

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. The no form of the command removes VAS interface type configuration.
Default	no vas-if-type
Parameters	<p>to-from-access — Used when two separate (to-from-access and to-from-network) interfaces are used for Value Added Service (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 the traffic is subject to regular routing but is not subject to AA divert, nor egress subscriber PBR.</p> <p>to-from-network — Used when two separate (to-from-access and to-from-network) interfaces are used for Value Added Service (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 regular routing can be applied.</p> <p>to-from-both — Used when a single interface is used for Value Added Service (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 the traffic is subject to regular routing but is not subject to AA divert, nor egress subscriber PBR.</p>

mode

Syntax	mode { strict loose strict-no-ecmp } no mode
Context	config>service>vprn>if>urfp-check config>service>vprn>nw-if>urfp-check
Description	This command specifies the mode of unicast RPF check. The no form of the command reverts to the default (strict) mode.
Default	strict
Parameters	<p>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.</p> <p>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.</p>

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.

lease-populate

Syntax	lease-populate [<i>nbr-of-leases</i>] lease-populate [<i>nbr-of-leases</i>] route-populate [pd] na [ta] lease-populate [<i>nbr-of-leases</i>] route-populate pd [na] [ta] [exclude] lease-populate [<i>nbr-of-leases</i>] route-populate [pd] [na] ta no lease-populate
Context	config>service>vprn>if>ipv6 config>service>vprn>if>ipv6>dhcp-relay
Description	<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"> • 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. <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"> • dhcpv6-pd • dhcpv6-na • dhcpv6-ta • dhcpv6-pd-excl
Parameters	<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>Values 1 — 8000</p> <p>route-populate —</p> <p>Values pd/na/ta — Create route based on specified option.</p> <p> exclude — Create blackhole route based on OPTION_PD_EXCLUDE.</p>

egress

Syntax	egress
Context	config>service>vprn>nw-if
Description	This command enables 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

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.
Default	none
Description	<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 l4-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 l4-load balancing is enabled) source port in the hash, ignore destination address/port.
 - destination** — Specifies using destination address and (if l4-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
- Note that 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	config>service>vprn>if
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	<p>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. Note that 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.</p> <p>interval <i>interval</i> — 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.</p> <p>Values 1— 6000 (Note that a zero value can be used by the SNMP agent to disable host-connectivity-verify.)</p> <p>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, etc.). 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.</p> <p>timeout <i>retry-timeout</i> — Specifies the timeout in seconds between consecutive retries of subscriber host connectivity verification checks, in case the host does not respond.</p> <p>Values 10— 60 seconds</p> <p>retry-count <i>count</i> — Specifies the number of retries that will be carried out before a subscriber host is considered to have failed the SHCV check.</p> <p>Values 2— 29</p> <p>family <i>family</i> — 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.</p> <p>Values 2— 29</p>

Interface ICMP Commands

icmp

Syntax	icmp
Context	config>service>vprn>if config>service>vprn>nw-if
Description	This command configures Internet Control Message Protocol (ICMP) parameters on a VPRN service.

mask-reply

Syntax	[no] mask-reply
Context	config>service>vprn>if>icmp config>service>vprn>nw-if>icmp
Description	<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 mask-reply 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 no form of this command disables replies to ICMP mask requests on the router interface.</p>
Default	mask-reply — Reply to ICMP mask requests.

redirects

Syntax	redirects [<i>number seconds</i>] no redirects
Context	config>service>vprn>if>icmp config>service>vprn>nw-if>icmp
Description	<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 redirects command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional <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>

Interface Commands

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of icmp redirects on the router interface.

Default **redirects 100 10** — Maximum of 100 redirect messages in 10 seconds.

Parameters *number* — The maximum number of ICMP redirect messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *seconds* of ICMP redirect messages that can be issued.

Values 1 — 60

tll-expired

Syntax **tll-expired** *number seconds*
no tll-expired

Context config>service>vprn>if>icmp
config>service>vprn>nw-if>icmp

Description Configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default tll-expired 100 10

Parameters *number* — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 — 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>vprn>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 *number* and *seconds* 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 per 10 second time interval.
- The **no** form of this command disables the generation of icmp destination unreachable messages on the router interface.
- Default** **unreachables 100 10**
- Parameters** *number* — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.
- Values** 10 — 1000
- seconds* — The time frame in seconds used to limit the *number* of ICMP unreachable messages that can be issued.
- Values** 1 — 60

Interface ICMP Commands

icmp

Syntax	icmp
Context	config>service>vprn>if config>service>vprn>nw-if
Description	This command configures Internet Control Message Protocol (ICMP) parameters on a VPRN service.

mask-reply

Syntax	[no] mask-reply
Context	config>service>vprn>if>icmp config>service>vprn>nw-if>icmp
Description	<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 mask-reply 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 no form of this command disables replies to ICMP mask requests on the router interface.</p>
Default	mask-reply — Reply to ICMP mask requests.

redirects

Syntax	redirects [<i>number seconds</i>] no redirects
Context	config>service>vprn>if>icmp config>service>vprn>nw-if>icmp
Description	<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 redirects command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional <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>

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of icmp redirects on the router interface.

Default **redirects 100 10** — Maximum of 100 redirect messages in 10 seconds.

Parameters *number* — The maximum number of ICMP redirect messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *seconds* of ICMP redirect messages that can be issued.

Values 1 — 60

tll-expired

Syntax **tll-expired** *number seconds*
no tll-expired

Context config>service>vprn>if>icmp
config>service>vprn>nw-if>icmp

Description Configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default tll-expired 100 10

Parameters *number* — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 — 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>vprn>if>icmp
config>service>vprn>nw-if>icmp

Interface Commands

- 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 *number* and *seconds* 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 per 10 second time interval.
- The **no** form of this command disables the generation of icmp destination unreachable messages on the router interface.
- Default** **unreachables 100 10**
- Parameters** *number* — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.
- Values** 10 — 1000
- seconds* — The time frame in seconds used to limit the *number* of ICMP unreachable messages that can be issued.
- Values** 1 — 60

Router Advertisement Commands

router-advertisement

Syntax	[no] router-advertisement
Context	config>service>vprn
Description	<p>This command configures router advertisement properties. By default, it is disabled for all IPv6 enabled interfaces.</p> <p>The no form of the command disables all IPv6 interface. However, the no interface <i>interface-name</i> command disables a specific interface.</p>
Default	disabled

dns-options

Syntax	[no] dns-options
Context	config>service>vprn>router-advertisement config>service>vprn>router-advertisement>interface
Description	<p>This command enables the context for configuration of DNS information for Stateless Address Auto-Configuration (SLAAC) hosts.</p> <p>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-advertisement>interface>dns-options>include-dns command.</p> <p>The no form of the command disables configuration of DNS information for Stateless Address Auto-Configuration (SLAAC) hosts.</p>
Default	disabled

dns-servers

Syntax	server <i>ipv6-address</i> no server
Context	config>service>vprn>router-advertisement>dns-options config>service>vprn>router-advertisement>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.
Default	none

Router Advertisement Commands

Parameters *ipv6-address* — Specify 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-advertisement>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 {seconds | infinite}**
no rdnss-lifetime

Context config>service>vprn>router-advertisement>dns-options
config>service>vprn>router-advertisement>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 infinite

Parameters **infinite** — specifies an infinite RDNSS lifetime.

seconds — Specifies the time in seconds.

Values 4— 3600

interface

Syntax **[no] interface *ip-int-name***

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 *ip-int-name* — Specify the interface name. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within 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	64
Parameters	<i>number</i> — Specifies the hop limit.
	Values 0 — 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	600
Parameters	<i>seconds</i> — Specifies the maximum interval in seconds between sending router advertisement messages.
	Values 4 — 1800

min-advertisement-interval

Syntax	[no] min-advertisement-interval <i>seconds</i>
Context	config>service>vprn>router-advert>if

Router Advertisement Commands

Description	This command configures the minimum interval between sending ICMPv6 neighbor discovery router advertisement messages.
Default	200
Parameters	<i>seconds</i> — Specify the minimum interval in seconds between sending ICMPv6 neighbor discovery router advertisement messages.
Values	3 — 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> — Specify the MTU for the nodes to use to send packets on the link.
Values	1280 — 9212

other-stateful-configuration

Syntax	[no] other-stateful-configuration
Description	This command sets the "Other configuration" flag. This flag indicates that DHCPv6lite is available for autoconfiguration of other (non-address) information such as DNS-related information or information on other servers in the network. 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.				
Default	none				
Parameters	<i>ip-prefix</i> — The IP prefix for prefix list entry in dotted decimal notation.				
Values	<table><tr><td>ipv4-prefix</td><td>a.b.c.d (host bits must be 0)</td></tr><tr><td>ipv4-prefix-length</td><td>0 — 32</td></tr></table>	ipv4-prefix	a.b.c.d (host bits must be 0)	ipv4-prefix-length	0 — 32
ipv4-prefix	a.b.c.d (host bits must be 0)				
ipv4-prefix-length	0 — 32				

ipv6-prefix	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d x: [0 — FFFF]H d: [0 — 255]D
ipv6-prefix-length	0 — 128

prefix-length — Specifies a route must match the most significant bits and have a prefix length.

Values 1 — 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 enabled

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 enabled

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 604800

Parameters *seconds* — Specifies the remaining length of time in seconds that this prefix will continue to be preferred.

infinite — 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> infinite }
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	2592000
Parameters	<i>seconds</i> — Specifies the remaining length of time in seconds that this prefix will continue to be valid. infinite — 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 — 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 — 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	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 — 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
	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.

Interface ICMP Commands

icmp

Syntax	icmp
Context	config>service>vprn>if config>service>vprn>nw-if
Description	This command configures Internet Control Message Protocol (ICMP) parameters on a VPRN service.

mask-reply

Syntax	[no] mask-reply
Context	config>service>vprn>if>icmp config>service>vprn>nw-if>icmp
Description	<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 mask-reply 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 no form of this command disables replies to ICMP mask requests on the router interface.</p>
Default	mask-reply — Reply to ICMP mask requests.

redirects

Syntax	redirects [<i>number seconds</i>] no redirects
Context	config>service>vprn>if>icmp config>service>vprn>nw-if>icmp
Description	<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 redirects command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional <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>

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of icmp redirects on the router interface.

Default **redirects 100 10** — Maximum of 100 redirect messages in 10 seconds.

Parameters *number* — The maximum number of ICMP redirect messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *seconds* of ICMP redirect messages that can be issued.

Values 1 — 60

ttl-expired

Syntax **ttl-expired** *number seconds*
no ttl-expired

Context config>service>vprn>if>icmp
config>service>vprn>nw-if>icmp

Description Configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default ttl-expired 100 10

Parameters *number* — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 — 60

if-attribute

Syntax **if-attribute**

Context 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 <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no admin-group <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no admin-group
Context	config>service>vprn>interface>if-attribute
Description	<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.</p> <p>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.</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>It should be noted that 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.</p> <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.</p>
Parameters	<i>group-name</i> — 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 <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no srlg-group <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no srlg-group
Context	config>service>vprn>interface>if-attribute
Description	<p>This command configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface.</p> <p>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.</p> <p>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.</p> <p>It should be noted that 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.</p> <p>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.</p>

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.

unreachables

Syntax **unreachables** [*number seconds*]
no unreachable

Context config>service>vprn>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 *number* and *seconds* 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 per 10 second time interval.

The **no** form of this command disables the generation of icmp destination unreachable messages on the router interface.

Default **unreachables 100 10**

Parameters *number* — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP unreachable messages that can be issued.

Values 1 — 60

ip-mtu

Syntax **ip-mtu** *octets*
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:

Router Advertisement Commands

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	octets —
Values	512 — 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-ig[:encap-val]</i> — Specifies the LAG ID.
Values	lag-id 1 — 800 encap-val 0 for null 0 — 4094 for dot1q

Interface SAP Commands

sap

Syntax	sap <i>sap-id</i> [create] no sap <i>sap-id</i>
Context	config>service>vprn>if
Description	<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 create 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 config interface <i>port-type port-id mode access</i> command. 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>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.</p>
Default	No SAPs are defined.
Special Cases	<p>VPRN — A VPRN SAP must be defined on an Ethernet interface.</p> <p>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.</p> <p>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 — 4094.</p>
Parameters	<p><i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition.</p> <p><i>port-id</i> — Specifies the physical port ID in the <i>slot/mda/port</i> format.</p> <p>If the card in the slot has Media Dependent Adapters (MDAs) installed, the <i>port-id</i> must be in the <i>slot_number/MDA_number/port_number</i> format. For example /2/3 specifies port 3 on MDA 2 in slot .</p>

Router Advertisement Commands

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.

create — Keyword used to create 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	<p>This command associates an aarp instance to a multi-homed SAP or spoke-sdp. This instance is paired with the same aarp-id in the same node or 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.</p> <p>The type specifies the role of this service point in the AARP: primary (dual-homed), secondary (dual-homed-secondary). The AA service attributes (app-profile, 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.</p>
Default	no aarp
Parameters	<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>Values 1 —</p> <p>type — Specifies the role of the SAP referenced by the AARP instance identified by AARP ID.</p> <p>Values dual-homed — the primary dual homed aa-subscriber side service point of an aarp instance, only supported for 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 IES and VPRN SAP and spoke-sdp.</p>

tod-suite

Syntax	tod-suite <i>tod-suite-name</i> no tod-suite
Context	config>service>vprn>if>sap

Description	This command applies a time-based policy (filter or QoS policy) to the SAP. The suite name must already exist in the config>system>cron context.
Default	no tod-suite
Parameters	<i>tod-suite-name</i> — Specifies collection of policies (ACLs, QoS) including time-ranges that define the full or partial behavior of a SAP or a subscriber. The suite can be applied to more than one SAP.

transit-policy

Syntax	transit-policy <i>ip-aasub-policy-id</i> transit-policy prefix <i>prefix-aasub-policy-id</i> no transit-ip-policy
Context	config>service>vprn>if>sap> config>service>vprn>if>spoke-sdp>
Description	This command associates a transit aa subscriber IP 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 the use of this policy assigned to services, 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-ip-policy <i>ip-aasub-policy-id</i> — An integer that identifies a transit IP profile entry. Values 1 — 65535 <i>prefix-aasub-policy-id</i> — An integer that identifies a prefix aasub-policy ID. Values 1 — 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 context.

Router Advertisement Commands

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

Default	Default accounting policy.
Parameters	<i>acct-policy-id</i> — Enter the accounting <i>policy-id</i> as configured in the config>log>accounting-policy context.
Values	1 — 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	bfd-enable no bfd-enable
Context	config>service>vprn>interface>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>interface>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	<p>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.</p> <p>When the no collect-stats command is issued the statistics are still accumulated by the cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent collect-stats command is issued then the counters written to the billing file include all the traffic while the no collect-stats command was in effect.</p>
Default	no collect-stats

source

Syntax	source ip-address
Context	config>service>ies>interface>sap>ip-tunnel
Description	<p>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.</p>
Default	no source
Parameters	<i>ip-address</i> — An IPv4 address or an IPv6 address.

remote-ip

Syntax	remote-ip ip-address no remote-ip
Context	config>service>ies>interface>sap>ip-tunnel
Description	<p>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.</p>

Router Advertisement Commands

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>interface>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 remote-ip
Parameters	<i>ip-address</i> — An IPv4 address or an IPv6 address.

Interface 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	<p>This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the interface.</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 (ip, mac, ip-mac, nh-mac) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.</p> <p>The no form of the command disables anti-spoof filtering on the SAP.</p>
Default	<p>Filter type default types:</p> <ul style="list-style-type: none"> • Non-Ethernet encapsulation default anti-spoof filter type — When enabled on a non-Ethernet encapsulated SAP, the anti-spoof filter default type is ip. • Ethernet encapsulated default anti-spoof filter type — When enabled on an Ethernet encapsulated SAP, the anti-spoof default type is ip-mac. • Default anti-spoof filter state — Anti-spoof filtering is disabled by default on the SAP.
Parameters	<p>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 will fail.</p> <p>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 will fail. The anti-spoof type mac command will also fail if the SAP does not support Ethernet encapsulation.</p> <p>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 will fail. 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.</p> <p>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.</p>

arp-limit

Syntax **arp-limit** *limit* [**log-only**] [**threshold percent**]

Router Advertisement Commands

	no arp-limit
Context	config>service>vprn>interface
Description	<p>This command configures the maximum amount of dynamic IPv4 ARP entries that can be learned on an IP interface.</p> <p>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.</p> <p>The no form of the command removes the arp-limit.</p>
Default	90 percent
Parameters	<p>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.</p> <p><i>percent</i> — The threshold value (as a percentage) that triggers a warning message to be sent.</p> <p>Values 0 — 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>Values 0 — 524288</p>

arp-populate

Syntax	[no] arp-populate
Context	config>service>vprn>if
Description	<p>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.</p> <p>The arp-populate command will fail if an existing static subscriber host on the SAP does not have both MAC and IP addresses specified.</p> <p>Once arp-populate is enabled, creating a static subscriber host on the SAP without both an IP address and MAC address will fail.</p> <p>arp-populate can only be enabled on VPRN interfaces supporting Ethernet encapsulation.</p> <p>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</p>

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 not enabled

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 seconds.

Default 5 seconds

Parameters *timer-multiple* — specifies the multiple of 100 ms that the ARP retry timer will be configured as.

Values 1 — 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

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 14400 seconds

Parameters *seconds* — The minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 — 65535

authentication-policy

Syntax	authentication-policy <i>name</i> no authentication-policy
Context	config>service>vprn>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, etc.), the entire string must be enclosed within double quotes.

delayed-enable

Syntax	delayed-enable <i>seconds</i> [init-only] no delayed-enable
Context	config>service>vprn>nw-if
Description	This command delays making interface operational by the specified number of seconds. In environments with many subscribers, it can take time to synchronize the subscriber state between peers when the subscriber-interface is enabled (perhaps, after a reboot). To ensure that the state has time to be synchronized, the delayed-enable timer can be specified. The optional parameter init-only can be added to use this timer only after a reboot.
Default	no delayed-enable
Parameters	<i>seconds</i> — Specifies the number of seconds to delay before the interface is operational. Values 1 — 1200 init-only — Delays the initialization of the subscriber-interface to give the rest of the system time to complete necessary tasks such as allowing routing protocols to converge and/or to allow MCS to sync the subscriber information. The delay only occurs immediately after a reboot.

calling-station-id

Syntax	calling-station-id <i>calling-station-id</i> no calling-station-id
Context	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	This attribute is not sent by default.

scheduling-class

Syntax	scheduling-class <i>class-id</i>
Context	config>service>vprn>if>sap
Description	This command specifies the scheduling class to use for this SAP.
Parameters	<i>class-id</i> — Specifies the scheduling class to use for this SAP.
	Values 0 — 3
	Default 0

flowspec

Syntax	flowspec no flowspec
Context	config>service>vprn>interface>sap>ingress config>service>vprn>interface>spoke-sdp>ingress config>service>ies>interface>sap>ingress config>service>ies>interface>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: <ol style="list-style-type: none"> 1. user-defined IPv4 filter entries 2. flowspec-derived filter entries 3. user-defined IPv4 filter default-action <p>The no form of the command removes IPv4 flowspec filtering from an IP interface.</p>
Default	No access interfaces have IPv4 flowspec enabled.

flowspec-ipv6

Syntax	flowspec-ipv6 no flowspec-ipv6
Context	config>service>vprn>interface>sap>ingress config>service>vprn>interface>spoke-sdp>ingress config>service>ies>interface>sap>ingress config>service>ies>interface>spoke-sdp>ingress
Description	This command enables IPv6 flowspec filtering on an access IP interface associated with a VPRN or IES service. Filtering is based on all of the IPv6 flowspec routes that have been received and accepted by the corresponding BGP instance. Ingress IPv6 traffic on an interface can be filtered by both a user-defined IPv6 filter and flowspec. Evaluation proceeds in this order:

Router Advertisement Commands

1. user-defined IPv6 filter entries
2. flowspec-derived filter entries
3. user-defined IPv6 filter default-action

The **no** form of the command removes IPv6 flowspec filtering from an IP interface.

Default No access interfaces have IPv6 flowspec enabled.

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
This command administratively enables host creation on this SAP.

Interface SAP Filter and QoS Policy Commands

egress

Syntax	egress
Context	config>service>vprn>if>sap
Description	<p>This command enables the context to configure egress SAP Quality of Service (QoS) policies and filter policies.</p> <p>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.</p>

ingress

Syntax	ingress
Context	config>service>vprn>if>sap config>service>vprn>if>sap
Description	<p>This command enables the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.</p> <p>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.</p>

agg-rate

Syntax	[no] agg-rate
Context	config>service>vprn>interface>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

Syntax	rate {max rate} no rate
Context	config>service>vprn>interface>sap>egress>agg-rate
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.

Router Advertisement Commands

(SAP, subscriber, VPORT etc.).

limit-unused-bandwidth

Syntax	[no] limit-unused-bandwidth
Context	config>service>vprn>interface>sap>egress>agg-rate
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>interface>sap>egress>agg-rate
Description	This command is used to enable (or disable) frame based accounting on all queues associated with the agg-rate context. Only supported on Ethernet ports. Not supported on HSMDA Ethernet ports.

agg-rate-limit

Syntax	agg-rate-limit <i>agg-rate</i> [queue-frame-based-accounting] no agg-rate-limit
Context	
Description	<p>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 will fail. If the agg-rate-limit command is specified, an attempt to bind a scheduler-policy to the SAP or multi-service site will fail.</p> <p>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 will fail. Once an agg-rate-limit has been assigned to a multi-service site, the scope cannot be changed to card level.</p> <p>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.</p> <p>If a port scheduler is not defined on the egress port, the queues are allowed to operate based on their own bandwidth parameters.</p> <p>The no form of the command removes the aggregate rate limit from the SAP or multi-service site.</p>

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 — 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** *ip-filter-id*
no filter

Context config>service>vprn>if>sap>egress
config>service>vprn>if>sap>ingress

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

Values 1 — 65535

flowspec

[no] flowspec

Context config>service>vprn>interface>sap>ingress
config>service>vprn>interface>spoke-sdp>ingress
config>service>vprn>network-interface>ingress

Router Advertisement Commands

Description This command enables flowspec filtering on an IP interface of a VPRN. Filtering is based on all of the flowspec routes that have been received and accepted by the VPRN. Ingress traffic on an IP interface can be filtered by both a user-defined ip filter and flowspec. Evaluation proceeds in this order:

1. user-defined ip filter entries with entry numbers less than the configured insert-point
2. flowspec-derived filter entries
3. user-defined ip filter entries with entry numbers greater than or equal to the configured insert-point
4. ip-filter default-action

The **no** form of the command removes flowspec filtering from an IP interface.

Default No interfaces have flowspec enabled.

flowspec-ipv6

Syntax **[no] flowspec-ipv6**

Context config>service>vprn>interface>sap>ingress
config>service>vprn>interface>spoke-sdp>ingress

Description This command enables flowspec filtering on an IP interface of the base router. Filtering is based on all of the flowspec routes that have been received and accepted by the base router. Ingress traffic on an IP interface can be filtered by both a user-defined ip filter and flowspec. In this case, the user-defined ip filter entries are evaluated before the flowspec routes and the default action of the user-defined ip filter applies as the very last rule.

The **no** form of the command removes flowspec filtering from an IP interface.

Default No interfaces have flowspec enabled.

wrr-weight

Syntax **wrr-weight** *value*
no wrr-weight

Context config>service>vprn>if>sap>egress>hsmda-queue-overider>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 *percentage* — Specifies the weight for the HSMDA queue.

Values 1— 32

wrr-policy

Syntax	wrr-policy <i>hsmda-wrr-policy-name</i> no wrr-policy
Context	config>service>vprn>if>sap>egress>hsmda-queue-overider
Description	This command associates an existing HSM DA weighted-round-robin (WRR) scheduling loop policy to the HSM DA queue.
Parameters	<i>hsmda-wrr-policy-name</i> — Specifies the existing HSM DA WRR policy name to associate to the queue.

secondary-shaper

	secondary-shaper <i>secondary-shaper-name</i> no secondary-shaper
Context	config>service>vprn>if>sap>egress>hsmda-queue-overider
Description	This command configures an HSM DA secondary shaper. Note that an shaper override can only be configured on an HSM DA SAP.
Parameters	<i>secondary-shaper-name</i> — Specifies a secondary shaper name up to 32 characters in length.

match-qinq-dot1p

Syntax	match-qinq-dot1p { top bottom } no match-qinq-dot1p
Context	config>service>vprn>if>sap>ingress
Description	<p>This command specifies which Dot1Q tag position Dot1P bits in a QinQ encapsulated packet should be used to evaluate Dot1P QoS classification.</p> <p>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.</p> <p>The no form of the command restores the default dot1p evaluation behavior for the SAP.</p> <p>By default, the bottom most service delineating Dot1Q tags Dot1P bits are used. The following table defines the default behavior for Dot1P evaluation when the match-qinq-dot1p command is not executed.</p>

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits

Port / SAP Type	Existing Packet Tags	PBits Used for Match
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	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. The following table defines the dot1p evaluation behavior when the top parameter is specified.

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 14: 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

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
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
Context	config>service>vprn>if>sap>egress config>service>vprn>sub-if>grp-if>sap>engress config>service>vprn>ipsec-if>sap>egress
Description	<p>This command associates a Quality of Service (QoS) policy with an ingress or 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 or IP interface. If the policy- id does not exist, an error will be returned.</p> <p>The <code>qos</code> command is used to associate both ingress and egress QoS policies. The <code>qos</code> 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> <p>The no form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.</p>
Default	none

Parameters	<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.</p> <p>1 — 65535</p> <p>port-redirect-group — 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-id 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 config>port>ethernet>access>egress.</p> <p>instance <i>instance-id</i> — Specifies the instance of the named egress port queue group on the IOM/IMM/XMA.</p> <p>Values 1 — 40960</p> <p>Default 1</p>
-------------------	--

qos

Syntax	<p>qos <i>policy-id</i> [shared-queuing multipoint-shared] fp-redirect-group <i>queue-group-name</i> instance <i>instance-id</i> no qos</p>
Context	<p>config>service>vprn>if>sap>ingress config>service>vprn>sub-if>grp-if>sap>ingress config>service>vprn>ipsec-if>sap>ingress</p>
Description	<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> <p>The no form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.</p>
Default	<p>none</p>

Router Advertisement Commands

Parameters *policy-id* — The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.

Values 1 — 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 — 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.

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 — This keyword 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 will fail.

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** *scheduler-policy-name*
no scheduler-policy

Context config>service>vprn>if>sap>ingress
config>service>vprn>if>sap>egress

Description 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 **config>qos>scheduler-policy** *scheduler-policy-name* 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 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 **no scheduler-policy** command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.

scheduler-policy-name: — The *scheduler-policy-name* 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.

lag-link-map-profile

Syntax	lag-link-map-profile <i>lnk-map-profile-id</i> no lag-link-map-profile
Context	config>service>vprn>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
Description	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. 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-

Router Advertisement Commands

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.

Default None — Each customer site must be explicitly created.

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 queues on the SAPs to no longer be orphaned. Existing schedulers on the site may cease to exist, causing 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, etc.), the entire string must be enclosed within double quotes.

static-host

Syntax **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*

Context config>service>vprn>if>sap

Description This command configures a static host on this SAP.

Syntax **ip** *ip-address* — Specifies the IPv4 unicast address.

mac *ieee-address* — Specify 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.

sla-profile *sla-profile-name*

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
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> no app-profile
Context	config>service>vprn>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.

inter-dest-id

Syntax	inter-dest-id <i>intermediate-destination-id</i> no inter-dest-id
Context	config>service>vprn>if>sap>static-host
Description	This command specifies to which intermediate destination (for example a DSLAM) this host belongs.
Parameters	<i>intermediate-destination-id</i> — Specifies the intermediate destination ID.

sla-profile

Syntax	sla-profile <i>sla-profile-name</i> no sla-profile
Context	config>service>vprn>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.

Router Advertisement Commands

Parameters *sla-profile-name* — Specifies the SLA profile name.

sub-profile

Syntax **sub-profile** *sub-profile-name*
no sub-profile

Context config>service>vprn>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>vprn>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>vprn>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-override

Syntax [**no**] **queue-override**

Context config>service>vprn>if>sap>egress
config>service>vprn>if>sap>ingress

Description This command enables 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.

queue

Syntax	[no] queue <i>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 — 32

adaptation-rule

Syntax	adaptation-rule [pir <i>adaptation-rule</i>] [cir <i>adaptation-rule</i>] no adaptation-rule
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. 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 <i>queue-id</i> 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 <i>queue-id</i> 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. <i>adaptation-rule</i> — 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 <i>percent</i> no avg-frame-overhead
Context	config>service>vprn>if>sap>egress>queue-override>queue
Description	<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> <ul style="list-style-type: none">• 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 x 0.1 or 1000 octets. <p>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 x 20 or 1000 octets.</p> <ul style="list-style-type: none">• 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 7500×1.1 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	<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.
Values	0 — 100

Router Advertisement Commands

cbs

Syntax	cbs <i>size-in-kbytes</i> no cbs
Context	config>service>vprn>if>sap>egress>queue-override>queue config>service>vprn>if>sap>ingress>queue-override>queue
Description	<p>This command can be used to override specific attributes of the specified queue's CBS parameters.</p> <p>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> <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.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.</p> <p>The no form of this command returns the CBS size to the default value.</p>
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. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).
Values	0 — 131072 or default

high-prio-only

Syntax	high-prio-only <i>percent</i> no high-prio-only
Context	config>service>vprn>if>sap>egress>queue-override>queue config>service>vprn>if>sap>ingress>queue-override>queue
Description	<p>This command can be used to override specific attributes of the specified queue's high-prio-only parameters. The high-prio-only command configures the percentage of buffer space for the queue, used exclusively by high priority packets.</p> <p>The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The high-prio-only parameter is used to override the default value derived from the network-queue command.</p> <p>The defined high-prio-only value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail</p>

execution. Attempting to set the **high-prio-only** value larger than the current MBS size will also result in an error and fail execution.

The **no** form of this command restores the default high priority reserved size.

Parameters *percent* — The *percent* parameter is the percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that none of the MBS of the queue will be reserved for high priority traffic. This does not affect RED slope operation for packets attempting to be queued.

Values 0 — 100 | default

mbs

Syntax **mbs** {*size-in-kbytes* | **default**}
no mbs

Context config>service>vprn>if>sap>egress>queue-override>queue
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 default

Parameters *size-in-kbytes* — The *size* parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.

For **sap>egress>queue-override>queue**:

Values 0 — 1073741824 or **default** in **bytes** or **kilobites**.

For **sap>egress>hsmda-queue-override>queue**:

Values [0 — 2625][**kilobytes**] | [0 — 2688000]**bytes** | **default**

mbs

Syntax	mbs { <i>size-in-kbytes</i> default } no mbs
Context	config>service>vprn>if>sap>ingress>queue-override>queue
Description	<p>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.</p> <p>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.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.</p> <p>The defined high-prio-only value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the high-prio-only value larger than the current MBS size will also result in an error and fail execution.</p> <p>The no form of this command returns the MBS size assigned to the queue to the value.</p>
Default	default
Parameters	<p><i>size-in-kbytes</i> — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.</p> <p>Values 0 — 131072 or default</p>

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	<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. 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 packets are preferentially queued by the system at egress and at</p>

subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile 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 *cir-level* and *cir-weight* 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 *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default **rate max cir 0** — The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir-rate** value.

Parameters *pir-rate* — 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 — 2000000000, **max**

Default max

cir-rate — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be 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 — 2000000000, **max**

Default 0

rate

Syntax **rate** *pir-rate*
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 *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR parameters (**max**, 0).

Router Advertisement Commands

Default *pir-rate* — 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 — 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** *scheduler-name*
no scheduler *scheduler-name*

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 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 queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier, except for schedulers created in tier 1). 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:

1. The maximum number of schedulers has not been configured.
2. The provided *scheduler-name* is valid.
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, etc.), 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.

parent

Syntax	parent [weight <i>weight</i>] [cir-weight <i>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	<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 no form of the command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.</p>
Default	no parent

Router Advertisement Commands

Parameters	<p>weight <i>weight</i> — 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.</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>Values 0 to 100</p> <p>Default 1</p> <p>cir-weight <i>cir-weight</i> — 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 queue or scheduler that the cir-weight pertains to has not reached the CIR and still has packets to transmit.</p> <p>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.</p> <p>Values 0 — 100</p> <p>Default 1</p>
-------------------	---

rate

Syntax	rate <i>pir-rate</i> [cir <i>cir-rate</i>] no rate
Context	config>service>vprn>if>sap>egress>sched-override>scheduler
Description	<p>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.</p> <p>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.</p> <p>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.</p>

The **no** form of this command returns all queues created with this *queue-id* by association with the QoS policy to the default PIR and CIR parameters.

Parameters

pir-rate — The **pir** parameter accepts a step multiplier value that specifies the multiplier used to determine the PIR rate at which the queue will operate. A value of 0 to 100000000 or the keyword **max** or **sum** is accepted. Any other value will result in an error without modifying the current PIR rate.

To calculate the actual PIR rate, the rate described by the queue's **rate** is multiplied by the *pir-rate*.

The SAP ingress context for PIR is independent of the defined forwarding class (fc) for the queue. The default **pir** and definable range is identical for each class. The PIR in effect for a queue defines the maximum rate at which the queue will be allowed to forward packets in a given second, thus shaping the queue's output.

The PIR parameter for SAP ingress queues do not have a negate (**no**) function. To return the queue's PIR rate to the default value, that value must be specified as the PIR value.

Values 1 — 100000000, **max**

Default max

cir cir-rate — The **cir** 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 0 to 250 or the keyword **max** is accepted. Any other value will result in an error without modifying the current CIR rate.

To calculate the actual CIR rate, the rate described by the **rate pir pir-rate** is multiplied by the *cir cir-rate*. If the **cir** is set to **max**, then the CIR rate is set to infinity.

The SAP ingress context for CIR is dependent on the defined forwarding class (fc) for the queue. The default CIR and definable range is different for each class. The CIR in effect for a queue defines both its profile (in or out) marking level as well as the relative importance compared to other queues for scheduling purposes during congestion periods.

Values 0 — 10000000, **max**, **sum**

Default sum

ETH-CFM Service Commands

eth-cfm

Syntax	eth-cfm
Context	config>service>vprn config>service>vprn>if>sap config>service>vprn>if>spoke-sdp
Description	This command enables the context to configure ETH-CFM parameters.

collect-lmm-stats

Syntax	collect-lmm-stats no collect-lmm-stats
Context	config>service>vprn>interface>sap>eth-cfm config>service>vprn>interface>spoke-sdp>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

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
Description	This command configures the ETH-CFM maintenance endpoint (MEP).
Parameters	<i>mep-id</i> — Specifies the maintenance association end point identifier. Values 1 — 8191 <i>md-index</i> — Specifies the maintenance domain (MD) index value. Values 1 — 4294967295

ma-index — Specifies the MA index value.

Values 1 — 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.

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
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.
Default	no interface-support-enabled (AIS will not be generated or stopped based on the state of the entity on) which the DOWN MEP is configured.

ccm-enable

Syntax	[no] ccm-enable
Context	config>service>vprn>if>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm>mep
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 <i>priority</i> no ccm-ltm-priority
Context	config>service>vprn>if>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm>mep
Description	This command specifies the priority value for CCMs and LTMs transmitted by the MEP. The no form of the command removes the priority value from the configuration.
Default	The highest priority on the bridge-port.
Parameters	<i>priority</i> — Specifies the priority of CCM and LTM messages. Values 0 — 7

ccm-padding-size

Syntax	[no] ccm-padding-size <i>ccm-padding</i>
Context	config>service>vprn>interface>sap>eth-cfm>mep config>service>vprn>interface>spoke-sdp>eth-cfm>mep config>service>vprn>subscriber-interface>group-interface>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.
Default	ccm-padding-size
Parameters	<i>ccm-padding</i> — specifies the byte size of the Optional Data TLV Values 3 — 1500

csf-enable

Syntax	[no] csf-enable
Context	config>service>vprn>interface>sap>eth-cfm>mep config>service>vprn>interface>spoke-sdp>eth-cfm>mep config>service>vprn>subscriber-interface>group-interface>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>interface>sap>eth-cfm>mep>cfs-enable config>service>vprn>interface>spoke-sdp>eth-cfm>mep>cfs-enable config>service>vprn>subscriber-interface>group-interface>sap>eth-cfm>mep>cfs-enable
Description	This command enables the multiplication factor applied to the receive time used to clear the CSF condition in increments of .5.
Default	3.5
Parameters	<i>multiplier-value</i> — Specifies the multiplier used for timing out CSF.
	Values 0.0, 2.0 .. 30.0

eth-test-enable

Syntax	[no] eth-test-enable
Context	config>service>vprn>if>sap>eth-cfm>mep
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: oam eth-cfm eth-test <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [<i>priority priority</i>] [<i>data-length data-length</i>] 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
Description	This command configures the test pattern for eth-test frames. The no form of the command removes the values from the configuration.
Parameters	all-zeros — Specifies to use all zeros in the test pattern. all-ones — Specifies to use all ones in the test pattern. crc-enable — Generates a CRC checksum.

ETH-CFM Service Commands

Default all-zeros

bit-error-threshold

Syntax **bit-error-threshold** *bit-errors*

Context config>service>vprn>if>sap>eth-cfm>mep
config>service>vprn>if>spoke-sdp>eth-cfm>mep

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

Default 1

Parameters *bit-errors* — Specifies the lowest priority defect.

Values 0 — 11840

one-way-delay-threshold

Syntax **one-way-delay-threshold** *time*

Context config>service>vprn>if>sap>eth-cfm
config>service>vprn>if>spoke-sdp>eth-cfm

Description This command enables one way delay threshold time limit.

Default 3 seconds

Parameters *priority* — Specifies the value for the threshold.

Values 0 — 600

snelch-ingress-levels

Syntax **snelch-ingress-levels** [*md-level* [*md-level*...]]
no snelch-ingress-levels

Context config>service>vprn>if>sap>eth-cfm
config>service>vprn>if>spoke-sdp>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..7]

tunnel-fault

Syntax	tunnel-fault {accept ignore}
Context	config>service>vprn>eth-cfm config>service>vprn>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 command ais-enable under epipe>sap>eth-cfm>ais-enable for more details. 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.
Parameters	accept — Share fate with the facility tunnel MEP ignore — Do not share fate with the facility tunnel MEP
Default	ignore (Service Level) accept (SAP Level for Epipe and VPLS)

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
Description	This command configures the fault propagation for the MEP.
Parameters	use-if-tlv — Specifies to use the interface TLV. suspend-ccm — Specifies to suspend the continuity check messages.

low-priority-defect

Syntax	low-priority-defect {allDef macRemErrXcon remErrXcon errXcon xcon noXcon}		
Context	config>service>vprn>if>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm>mep		
Description	This command specifies the lowest priority defect that is allowed to generate a fault alarm.		
Default	macRemErrXcon		
Values	allDef		DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
	macRemErrXcon		Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
	remErrXcon		Only DefRemoteCCM, DefErrorCCM, and DefXconCCM
	errXcon		Only DefErrorCCM and DefXconCCM
	xcon		Only DefXconCCM; or
	noXcon		No defects DefXcon or lower are to be reported

Interface VRRP Commands

vrrp

Syntax	vrrp <i>virtual-router-id</i> [owner] no vrrp <i>virtual-router-id</i>
Context	config>service>vprn>if
Description	<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 owner. The nodal context of <code>vrrp virtual-router-id</code> is used to define the configuration parameters for the VRID.</p> <p>The no form of this command removes the specified VRID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the VRID. The VRID does not need to be shutdown in order to remove the virtual router instance.</p>
Default	No default
Parameters	<p><i>virtual-router-id</i> — The <i>virtual-router-id</i> parameter specifies a new virtual router ID or one that can be modified on the IP interface.</p> <p>Values 1 — 255</p>

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	<p>The authentication-key command, within the <code>vrrp virtual-router-id</code> context, is used to assign a simple text password authentication key to generate master VRRP advertisement messages and validate received VRRP advertisement messages.</p> <p>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.</p> <p>To change the current in-use password key on multiple virtual router instances:</p> <ul style="list-style-type: none"> • 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 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** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

backup

Syntax **[no] backup ip-address**

Context config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp

Description This command configures virtual router IP addresses for the interface.

bfd-enable

Syntax	bfd-enable interface <i>interface-name</i> dst-ip <i>ip-address</i> bfd-enable service-id 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> no bfd-enable service-id interface <i>interface-name</i> dst-ip <i>ip-address</i>
Context	config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp
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 300ms. 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.
Default	none
Parameters	<i>service-id</i> — Specifies the service ID of the interface running BFD. Values <i>service-id</i> : 1 — 2147483648 <i>svc-name</i> : Specifies an existing service name up to 64 characters in length. No service ID indicates a network interface. interface <i>interface-name</i> — Specifies the name of the interface running BFD. dst-ip <i>ip-address</i> — Specifies the destination address to be used for the BFD session.

init-delay

Syntax	init-delay <i>seconds</i> no init-delay
Context	config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp
Description	This command configures a VRRP initialization delay timer.
Default	no init-delay
Parameters	<i>seconds</i> — Specifies the initialization delay timer for VRRP, in seconds. Values 1 — 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	<i>seconds</i> — The number of seconds that will transpire before the advertisement timer expires.
Values	1 — 255
Default	1

milliseconds *milliseconds* — Specifies the milliseconds time interval between sending advertisement messages. This parameter is not supported on single-slot chassis.

Values 100 — 900

ping-reply

Syntax	[no] ping-reply
Context	config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp
Description	<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 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.</p> <p>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.</p>
Default	no ping-reply

policy

Syntax	policy vrrp-policy-id no policy
Context	config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp
Description	This command associates a VRRP priority control policy with the virtual router instance (non-owner context only).
Parameters	<i>vrrp-policy-id</i> — Specifies a VRRP priority control policy.
Values	1 — 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 — 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.

ETH-CFM Service Commands

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 **traceroute-reply** status.

Default no traceroute-reply

PIM Commands

pim

Syntax	[no] pim
Context	config>service>vprn
Description	<p>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).</p> <p>The no form of the command deletes the PIM protocol instance removing all associated configuration parameters.</p>
Default	none

apply-to

Syntax	apply-to {all none}
Context	config>service>vprn>pim
Description	<p>This command creates a PIM interface with default parameters.</p> <p>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.</p> <p>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.</p>
Default	none (keyword)
Parameters	<p>all — Specifies that all VPRN and non-VPRN interfaces are automatically applied in PIM.</p> <p>none — No interfaces are automatically applied in PIM. PIM interfaces must be manually configured.</p>

grt-extranet

Syntax	[no] grt-extranet
Context	config>service>vprn>pim
Description	This command enables 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...(upto 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.
Default	none
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>] <i>policy-name</i>] 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, etc.), the entire string must be enclosed within 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	<p>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.</p> <p>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.</p> <p>The shutdown command can be used to disable an interface without removing the configuration for the interface.</p>
Default	Interfaces are activated in PIM based on the apply-to command.
Parameters	<i>ip-int-name</i> — Specify the interface name. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

assert-period

Syntax	assert-period <i>assert-period</i> no assert-period
Context	config>service>vprn>pim>if
Description	<p>This command configures the period in seconds for periodic refreshes of PIM Assert messages on an interface.</p> <p>The no form of the command reverts to the default.</p>
Default	60
	<i>assert-period</i> — Specifies the period, in seconds, for periodic refreshes of PIM Assert messages on an interface.
Values	1 — 300

bfd-enable

Syntax	[no] bfd-enable [ipv4 ipv6]
Context	config>service>vprn>pim>if
Description	<p>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.</p> <p>The no form of this command removes BFD from the associated IGP protocol adjacency.</p>

PIM Commands

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 reverts to the default value.

Default 30

Parameters *hello-interval* — Specifies the hello interval in seconds. A 0 (zero) value disables the sending of hello messages.

Values 0 — 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.

The **hello-multiplier** in conjunction with the **hello-interval** determines the holdtime for a PIM neighbor.

Parameters *deci-units* — Specify the value, specified in multiples of 0.1, for the formula used to calculate the hello-holdtime based on the hello-multiplier:

$(\text{hello-interval} * \text{hello-multiplier}) / 10$

This allows the PIMv2 default timeout of 3.5 seconds to be supported.

Values 20 — 100

Default 35

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 enabled

instant-prune-echo

`[no] instant-prune-echo`

max-groups

Syntax `max-groups value`
`no max-groups`

Context `config>service>vprn>pim>if`

Description This command configures the maximum number of groups for which PIM can have 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 *value* — Specifies the maximum number of groups for this interface.

Values 1 — 16000

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).
Parameters	ipv4 — Enables dynamic mLDP 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 . ipv6 — Enables dynamic mLDP in-band signaling for IPv6 PIM joins. IPv6 multicast must be enabled; see ipv6-multicast-disable .
Default	no p2mp-ldp-tree-join

priority

Syntax	priority <i>dr-priority</i> 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	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 — 4294967295

sticky-dr

Syntax	sticky-dr [priority <i>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	disabled
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 — 4294967295

three-way-hello

Syntax	three-way-hello [compatibility-mode] 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 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

mc-ecmp-balance

Syntax	[no] mc-ecmp-balance
Context	config>service>vprn>pim
Description	This command enables multicast balancing of traffic over ECMP links. When enabled, each multicast stream that needs to be forwarded over an ECMP link will be re-evaluated for the total multicast bandwidth utilization. Re-evaluation occurs on the ECMP interface in question. The no form of the command disables the multicast balancing.

mc-ecmp-balance-hold

Syntax	mc-ecmp-balance-hold <i>minutes</i> 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.

non-dr-attract-traffic

Syntax	[no] non-dr-attract-traffic
Context	config>service>vprn>pim
Description	<p>This command specifies whether the router should ignore the designated router state and attract traffic even when it is not the designater router.</p> <p>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. Note that while using this flag the router may not send the stream down to the DSLAM while not DR.</p> <p>When enabled, the designated router state is ignored. When disabled, no non-dr-attract-traffic, the designated router value is honored.</p>
Default	no non-dr-attract-traffic

rp

Syntax	rp
Context	config>service>vprn>pim
Description	<p>This command enables access to the context to configure the rendezvous point (RP)) of a PIM protocol instance.</p> <p>An Alcatel-Lucent PIM router acting as an RP must respond to a PIM register message specifying an SSM multicast group address by sending to the first hop router stop register 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 of 232/8 or from a multicast group address range that was explicitly configured for SSM.</p>
Default	rp enabled when PIM is enabled.

anycast

Syntax	[no] anycast <i>rp-ip-address</i>
Context	config>service>vprn>pim>rp
Description	<p>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.</p> <p>The no form of the command removes the anycast instance from the configuration.</p>
Default	none
Parameters	<p><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.</p> <p>Values Any valid loopback address configured on the node.</p>

rp-set-peer

Syntax	[no] rp-set-peer <i>ip-address</i>
Context	config>service>vprn>pim>rp>anycast
Description	<p>This command configures a peer in the anycast rp-set. The address identifies the address used by the other node as the RP candidacy address for the same multicast group address range as configured on this node.</p> <p>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.</p> <p>Although there is no set maximum of addresses that can be configured in an rp-set, up to 15 multicast addresses is recommended.</p> <p>The no form of the command removes an entry from the list.</p>
Default	None
Parameters	<p><i>ip-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.</p>

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.
- Note that chassis mode D or higher must be enabled for **auto-rp-discovery**.
- The **no** form of the command disables auto RP.
- Default** disabled

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.
- Default** none
- Parameters** *policy-name* — Specify 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.
- Default** none
- Parameters** *policy-name* — Specify 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

PIM Commands

- Description** This command enables the context to configure a local rendezvous point (RP) of a PIM protocol instance.
- Default** Enabled when PIM is enabled.

address

- Syntax** `[no] address ip-address`
- Context** `config>service>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** *ip-address* — The static IP address of the RP. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.
- Values** 1.0.0.0 – 223.255.255.255

address

- Syntax** `[no] address ipv6-address`
- 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** *ipv6-address* — The static IP address of the RP. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.
- Values** `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`

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	30
Parameters	<i>hash-mask-length</i> — The hash mask length.
	Values 0 — 32

hash-mask-length

Syntax	hash-mask-length <i>hash-mask-length</i> no hash-mask-length
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	126
Parameters	<i>hash-mask-length</i> — The hash mask length.
	Values 0 — 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 — 255
	Default 0 (the router is not eligible to be the bootstrap router)

PIM Commands

ipv6

Syntax	ipv6
Context	config>service>vprn>pim>rp
Description	<p>This command enables access to the context to configure the rendezvous point (RP) of a PIM IPv6 protocol instance.</p> <p>An Alcatel-Lucent 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.</p>
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	<p>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.</p> <p>The no form of the command removes the anycast instance from the configuration.</p>
Default	none
Parameters	<p><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 level. —thi highlighted text is invalid, don't add here and please remove from IPv4 command</p> <p>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..FFFF]H d [0..255]D</p>

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.

Default None

Parameters *ipv6-address* — Specifies the address used by the other node as the RP candidacy address for the same multicast group address range as configured on this node.

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..FFFF]H
 d [0..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.

Default none

Parameters *ipv6-address* — Specify the addresses or address ranges that this router can be an RP.

prefix-length — Specify the address prefix length.

Values *ipv6-address* x:x:x:x:x:x:x (eight 16-bit pieces)
 x:x:x:x:x:d.d.d.d

PIM Commands

x [0..FFFF] H
d [0..255] D
prefix-length [8..128] //for embedded-rp
prefix-length [16..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.

Default none

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..FFFF] H
d [0..255] D
prefix-length [8..128]

rp-candidate

Syntax rp-candidate

Context config>service>vprn>pim>rp
config>service>vprn>pim>rp>ipv6

Description This command enables the context to configure the candidate rendezvous point (RP) parameters.

Default Enabled when PIM is enabled.

group-range

Syntax [no] group-range {*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.

Default none

Parameters *ip-prefix* — Specify the addresses or address ranges that this router can be an RP.

Values *ipv4-prefix* - a.b.c.d
 ipv4-prefix-le - [0..32]
 ipv6-prefix - x:x:x:x:x:x:x (eight 16-bitpieces)
 x:x:x:x:x:d.d.d.d
 x - [0..FFFF]H
 d - [0..255]D
 ipv6-prefix-le - [0..128]

mask — Specify the address mask with the address to define a range of addresses.

netmask — Specify the subnet mask in dotted decimal notation.

Values :a.b.c.d (network bits all 1 and host bits all 0)

holdtime

Syntax **holdtime** *holdtime*
no holdtime *holdtime*

Context config>service>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 150

Parameters *holdtime* — Specify the length of time, in seconds, that neighbor should consider the sending router to be operational.

Values 0 — 255

priority

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

Context config>router>pim>rp>local
 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.

PIM Commands

Use the **no** form of this command to revert to the default value.

Default 1

Parameters *priority* — Specify the priority to become the designated router. The higher the value the more likely the router will become the RP.

Values 0 — 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.

Default none

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.

Default none

group-prefix

Syntax [**no**] **group-prefix** {*grp-ip-address/mask* | *grp-ip-address netmask*}

Context config>service>vprn>pim>rp>static

Context 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.

Default none

Parameters *grp-ip-address* — Specify the multicast IP address.

mask — Defines the mask of the multicast-ip-address.

Values 4 — 32

netmask — Enter the subnet mask in dotted decimal notation.

Values 0.0.0.0 — 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	[no] rpf-table {rtable-m rtable-u both}
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 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: <ul style="list-style-type: none"> a) Use unicast route table only b) Use multicast route table only or c) Use both the route tables.
Default	rpf-table rtable-u
Parameters	<p>rtable-m — pecified that only the multicast route table is to be used by the multicast protocol (PIM) for IPv4 RPF checks. This route table contains routes submitted by static routes and OSPF.</p> <p>rtable-u — Specifies that only the unicast route table is to be used by the multicast protocol (PIM) for IPv4 RPF checks. This route table contains routes submitted by all the unicast routing protocols.</p> <p>both — Specifies that PIM always lookup first in the multicast route table, and if there is a route, PIM 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 rtable-u.</p>

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> — Specify the multicast group address. <i>mask</i> — Defines the mask of the multicast-ip-address. Values 4 — 32 <i>netmask</i> — Enter the subnet mask in dotted decimal notation. Values 0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0) <i>spt-threshold</i> — Specifies the configured threshold in kilo-bits per second(kbps) for the group to which this (S,G) belongs. For a group G configured with a threshold, switchover to SPT for an (S,G) is attempted only if the (S,G)'s rate exceeds this configured threshold.

ssm-assert-compatible-mode

Syntax	ssm-assert-compatible-mode [enable disable]
Context	config>service>vprn>pim
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	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	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.
Default	none

C-MLDP Commands

mld

Syntax	[no] mld
Context	config>service>vprn
Description	This command enables 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	<p>This command enables 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.</p> <p>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.</p>
Default	no interface — No interfaces are defined.
Parameters	<p><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, etc.), the entire string must be enclosed within double quotes.</p> <p>If the IP interface name does not exist or does not have an IP address configured an error message will be returned.</p> <p>If the IP interface exists in a different area it will be moved to this area.</p>

disable-router-alert-check

Syntax	[no] disable-router-alert-check
Context	config>service>vprn>mld>interface
Description	This command enables router alert checking for MLD messages received on this interface.

The no form of the command disables the router alert checking.

Default none

import

Syntax **import** *policy-name*
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** — No import policy specified.

Parameters *policy-name* — The route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. 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 — 16000

query-interval

Syntax	query-interval <i>seconds</i> 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	<i>seconds</i> — The time frequency, in seconds, that the router transmits general host-query messages. Values 2 — 1024

query-last-member-interval

Syntax	query-last-member-interval <i>seconds</i>
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 — 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 the length of time to wait to receive a response to the host-query message from the host. Values 1 — 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.
Default	none

group

Syntax	[no] group <i>ipv6-address</i>
Context	config>service>vprn>mld>interface>static
Description	<p>This command enables 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.</p> <p>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.</p> <p>The no form of the command removes the IPv6 address from the configuration.</p>
Default	none
Parameters	<i>ipv6-address</i> — Specifies an MLD multicast group address that receives data on an interface. The IP address must be unique for each static group.

source

Syntax	[no] source <i>ipv6-address</i>
Context	config>service>vprn>mld>interface>static>group
Description	<p>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 source(s) that the traffic is expected.</p> <p>The source command is mutually exclusive with the specification of individual sources for the same group.</p> <p>The source command, in combination with the group, is used to create a specific (S,G) static group entry.</p> <p>The no form of the command removes the source from the configuration.</p>
Default	none
Parameters	<i>ip-address</i> — 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.
Default	none

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	1
Parameters	<i>version</i> — 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	2
Parameters	<i>robust-count</i> — Specify the robust count value. Values 2 — 10

ssm-translate

Syntax	ssm-translate
Context	config>service>vprn>mld
Description	This command enables 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	<p><i>start</i> — An IP address that specifies the start of the group range.</p> <p><i>end</i> — 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 <i>start</i> value.</p>

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 <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.

Network Interface Commands

network-interface

- Syntax** `network-interface interface-name [create]`
`no network-interface interface-name`
- 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.

Network Ingress Commands

network

Syntax	network
Context	config>service>vprn
Description	This command enables the context to configure network parameters for the VPRN service.

ingress

Syntax	ingress
Context	config>service>vprn>network
Description	This command enables the context to configure network ingress parameters for the VPRN service.

qos

Syntax	qos network-policy-id fp-redirect-group queue-group-name instance instance-id 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 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:

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.
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 will fail 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 will use 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. Note that multicast traffic would use the ingress network queues or queue group related to the network interface.

Network Ingress Commands

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 *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

Values 1—65535

fp-redirect-group *queue-group-name* — Specifies the name of the queue group template up to 32 characters in length.

instance *instance-id* — Specifies the identification of a specific instance of the queue-group.

Values 1—65535

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 [cir rate] [pir rate] 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 affect 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 routes to other BGP peers. By default, BGP only advertises BGP routes to other BGP peers if a given BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the most preferred route within the system for a given destination.
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	<p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>The no form of the command used at the group level reverts to the value defined at the group level.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>
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	<p>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.</p>
Default	no always-compare-med
Parameters	<p>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.</p> <p>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.</p>

strike-as — Specifies BGP paths to be compared even with different neighbor AS.

as-path-ignore

Syntax	[no] as-path-ignore
Context	config>service>vprn>bgp
Description	This command determines whether the AS path is used to determine the best BGP route. If this option is present, the AS paths of incoming routes are not used in the route selection process. The no form of the command removes the parameter from the configuration.
Default	no as-path-ignore

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

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	as-override is not enabled by default.

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [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	Authentication is disabled and the authentication password is empty.
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 parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash parameter specified. hash2 — Specifies the key is entered in a more complex encrypted form. If the hash2 parameter is not used, the less encrypted hash form is assumed.

auth-keychain

Syntax	auth-keychain <i>name</i>
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
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.

backup-path

Syntax	[no] backup-path [ipv4] [ipv6]
Context	config>router config>router>bgp 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. In some cases prefix independent convergence may require use of FP2 or later IOMs/IMMs/XMAs. By default, IPv4 and IPv6 prefixes do not have a backup path installed in the IOM.
Default	no backup-path
Parameters	ipv4 — enable the use of a backup path for BGP-learned IPv4 prefixes. ipv6 — enable the use of a backup path for BGP-learned IPv6 prefixes.

best-path-selection

Syntax	best-path-selection
Context	config>service>vprn>bgp
Description	This command enables path selection configuration.

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 <i>cluster-id</i> no cluster
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	<p>This command configures the cluster ID for a route reflector server.</p> <p>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.</p> <p>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.</p> <p>For redundancy, a cluster can have multiple route reflectors.</p> <p>Confederations can also be used to remove the full IBGP mesh requirement within an AS.</p> <p>The no form of the command deletes the cluster ID and effectively disables the Route Reflection for the given group.</p>
Default	no cluster — No cluster ID is defined.
Parameters	<i>cluster-id</i> — The route reflector cluster ID is expressed in dot decimal notation.
	Values Any 32 bit number in dot decimal notation. (0.0.0.1 — 255.255.255.255)

connect-retry

Syntax	connect-retry <i>seconds</i> no connect-retry
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	<p>This command configures the BGP connect retry timer value in seconds.</p> <p>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.</p> <p>The no form of the command used at the global level reverts to the default value.</p> <p>The no form of the command used at the group level reverts to the value defined at the global level.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>
Default	120 seconds
Parameters	<i>seconds</i> — The BGP Connect Retry timer value in seconds, expressed as a decimal integer.

Values 1 — 65535

damp-peer-oscillations

Syntax	damp-peer-oscillations [idle-hold-time <i>initial-wait</i> second-wait <i>max-wait</i>] [error-interval <i>minutes</i>]
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	<p>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>.</p> <p>The default behavior, which applies when no damp-peer-oscillations is configured, is to immediately transition out of the idle-state after every reset.</p>
Default	<i>no damp-peer-oscillations</i>
Parameters	<p><i>initial-wait</i> — The amount of time, in minutes, that a session remains in the idle-state after it has been stable for a while.</p> <p>Values 0 — 2048</p> <p>Default 0</p> <p><i>second-wait</i> — A period of time, in minutes, that is doubled after each repeated session failure that occurs within a relatively short span of time.</p> <p>Values 0 — 2048</p> <p>Default 5</p> <p><i>max-wait</i> — The maximum amount of time, in minutes, that a session remains in the idle-state after it has experienced repeated instability.</p> <p>Values 0 — 2048</p> <p>Default 60</p> <p><i>minutes</i> — 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.</p> <p>Values 0 — 2048</p> <p>Default 30</p>

damping

Syntax	[no] damping
Context	config>service>vprn>bgp

```
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor
```

Description This command enables BGP 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 will strictly support 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 — Client routes are reflected to all client peers.

disable-communities

Syntax disable-communities [standard] [extended]
no disable-communities

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures BGP to disable sending communities.

Parameters **standard** — Specifies standard communities that existed before VPRNs or 2547.
extended — Specifies BGP communities used were expanded after the concept of 2547 was introduced, to include handling the VRF target.

disable-fast-external-failover

Syntax [no] disable-fast-external-failover

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures BGP fast external failover.

eibgp-loadbalance

Syntax	[no] eibgp-loadbalance
Context	config>service>vprn>bgp
Description	<p>This command enables eiBGP load sharing so routes with both MP-BGP and IPv4 next-hops can be used simultaneously.</p> <p>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.</p> <p>The no form of the command used at the global level reverts to default values.</p>
Default	no eibgp-loadbalance — Multipath disabled.

enable-bgp-vpn-backup

Syntax	enable-bgp-vpn-backup [ipv4] [ipv6] no enable-bgp-vpn-backup
Context	config>service>vprn
Description	<p>This command enables BGP-VPN routes imported into the VPRN to have backup paths calculated for them (when they are the best path) and to be considered as backup path candidates (for other VPN-IP routes and VPRN BGP routes learned from CEs).</p>
Default	no enable-bgp-vpn-backup

ebgp-link-bandwidth

Syntax	ebgp-link-bandwidth <i>family</i> [<i>family</i> ... (up to 2 max)] no ebgp-link-bandwidth				
Context	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor				
Description	<p>When the ebgp-link-bandwidth command is configured, BGP automatically adds a link-bandwidth extended community to every route (of the selected types) received from directly connected (single-hop) EBGP peers within the scope of the command.</p> <p>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.</p>				
Default	no ebgp-link-bandwidth — No link bandwidth extended community is automatically added to received BGP routes.				
Parameters	<i>family</i> — The BGP address family.				
Values	<table border="0"> <tr> <td><i>ipv4</i></td> <td>The command applies to IPv4 and label-IPv4 routes.</td> </tr> <tr> <td><i>ipv6</i></td> <td>The command applies to IPv6 and 6PE routes.</td> </tr> </table>	<i>ipv4</i>	The command applies to IPv4 and label-IPv4 routes.	<i>ipv6</i>	The command applies to IPv6 and 6PE routes.
<i>ipv4</i>	The command applies to IPv4 and label-IPv4 routes.				
<i>ipv6</i>	The command applies to IPv6 and 6PE routes.				

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 or disables graceful-restart for all VPRN BGP peers.

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

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> — The restart-time that is advertised in the router's graceful-restart capability. Values 0 — 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.
Parameters	<i>time</i> — [1..3600 seconds]
Default	360 seconds

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 fault-tolerance

export

Syntax	export <i>policy</i> [<i>policy</i>...] no export
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command specifies the export policies to be used to control routes advertised to BGP neighbors. When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be configured. The first policy that matches is applied. Note that if a non-existent route policy is applied to a VPRN instance, the CLI generates a warning message. This message is only generated at an interactive CLI session and the route policy association is made. No warning message is generated when a non-existent route policy is applied to a VPRN instance in a configuration file or when SNMP is used. The no form of this command removes all route policy names from the export list.
Default	no export — BGP advertises routes from other BGP routes but does not advertise any routes from other protocols unless directed by an export policy.
Parameters	<i>policy</i> — A route policy statement name.

family

Syntax	family [ipv4][mcast-ipv4] [flow-ipv6] [flow-ipv4] no family
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command specifies the address families to be negotiated with one or more multi-protocol BGP peers of the VPRN.

BGP Commands

The **no** form of the command removes the specified address family from the associated BGP sessions.

Default ipv4

Parameters *ipv4* — Provisions IPv4 support.

mcast-ipv4 — Provisions Multicast IPv4 support.

[flow-ipv6] — Exchanges IPv4 flowspec routes belonging to AFI 2 and SAFI 133.

[flow-ipv4] — Specifies to use an address of variable size consisting of 1 or 2-byte NLRI(Network Layer Reachability Information) length followed by a variable length NLRI value.

flowspec-validate

Syntax **flowspec-validate**
no flowspec-validate

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command enables/disables validation of received flowspec routes. A flow route with a destination prefix subcomponent received from a particular peer is considered valid if and only if that peer also advertised the best unicast route to the destination prefix and any of its more-specific components. If validation is enabled and a flowspec route is not valid, it is not eligible for import into the RIB, it is not used for filtering, and it is not propagated to other flowspec peers.

The **no** form of the command disables the validation procedure.

Default no flowspec-validate

group

Syntax **group name [esm-dynamic-peer]**
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 shutdown before it can be deleted.

Default None — No peer groups are defined.

Parameters *name* — 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, etc.), the entire string must be enclosed within double quotes.

esm-dynamic-peer — This flag designates that the given BGP group will be used by BGP peers created dynamically based on subscriber-hosts pointing to corresponding BGP peering policy.

There can 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 shutdown, the command will not result in any action except a warning message on the console indicating that neighbor is still administratively up.
Default	none — No neighbors are defined.
Parameters	<i>ip-address</i> — The IP address of the BGP peer router in dotted decimal notation.
Values	ipv4-address : a.b.c.d

family

Syntax	family [ipv4] [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. 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	<p>This command configures the BGP hold time, expressed in seconds.</p> <p>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.</p> <p>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:</p> <ol style="list-style-type: none"> 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. <p>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.</p>
Default	90 seconds
Parameters	<p><i>seconds</i> — The hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is up permanently.</p> <p>Values 0, 3 — 65535</p> <p><i>seconds2</i> — The minimum hold-time that will be accepted for the session. If the peer proposes a hold-time lower than this value the session attempt will be rejected.</p>

ibgp-multipath

Syntax	[no] ibgp-multipath
Context	config>service>vprn>bgp
Description	<p>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 nexthops.</p> <p>The no form of the command disables the IBGP multipath load balancing feature.</p>
Default	none

import

Syntax	import <i>policy</i> [<i>policy</i> ...] no import
Context	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	This command specifies the import policies to be used to control routes advertised to BGP neighbors. Route policies are configured in the config>router>policy-options context. When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be specified. The first policy that matches is applied. The no form of this command removes all route policy names from the import list.
Default	no import — BGP accepts all routes from configured BGP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.
Parameters	<i>policy</i> — A route policy statement name.

keepalive

Syntax	keepalive <i>seconds</i> 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 <i>seconds</i> 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 keepalive value, then the keepalive value is reset to one third of the specified hold-time interval. If the hold-time interval is set to 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

BGP Commands

Parameters *seconds* — The keepalive timer in seconds, expressed as a decimal integer.

Values 0 — 21845

local-address

Syntax **local-address** *ip-address*
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 router 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.

Default **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 will remove any virtual AS number configured. The **no** form of the command used at the group level reverts to the value defined at the global level. The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default	no local-as
Parameters	<i>as-number</i> — The virtual autonomous system number, expressed as a decimal integer.
Values	1 — 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	<p>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.</p> <p>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.</p> <p>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.</p> <p>The no form of the command used at the group level reverts to the value defined at the global level.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>

BGP Commands

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 — 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	<p>This command configures how the BGP peer session handles loop detection in the AS path.</p> <p>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.</p> <p>Note that dynamic configuration changes of loop-detect are not recognized.</p> <p>The no form of the command used at the global level reverts to default, which is loop-detect ignore-loop.</p> <p>The no form of the command used at the group level reverts to the value defined at the global level.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>
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	<p>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.</p> <p>The specified value can be overridden by any value set via a route policy.</p>

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 — 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— 255

multihop

Syntax	multihop <i>tvl-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.

BGP Commands

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 *ttl-value* — The TTL value, expressed as a decimal integer.

Values 1 — 255

multipath

Syntax **multipath** *max-paths* [**eibgp**]
no multipath

Context config>service>vprn>bgp

Description This command enables BGP multipath.

When multipath is enabled BGP load shares traffic across multiple links. Multipath can be configured to load share traffic across a maximum of 16 routes. If the equal cost routes available are more than the configured value, then routes with the lowest next-hop IP address value are chosen.

This configuration parameter 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.

Multipath is effectively disabled if the value is set to one. When multipath is disabled, and multiple equal cost routes are available, the route with the lowest next-hop IP address will be used.

The **no** form of the command used at the global level reverts to default values.

Default **no multipath** — Multipath disabled.

Parameters *integer* — The number of equal cost routes to use for multipath routing. If more equal cost routes exist than the configured value, routes with the lowest next-hop value are chosen. Setting this value to 1 disables multipath.

Values 1 — 16

eibgp — Enables EIBGP load balancing so that routes with both MP-BGP and IPv4 next-hops can be used simultaneously. Enabling this option will disable the nexthop type (MP-BGP or IPv4 and also the next-hop metric comparison).

next-hop-resolution

Syntax	next-hop-resolution
Context	config>service>vprn>bgp
Description	This command enables 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.
Default	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.
Default	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	<p>This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.</p> <p>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</p> <p>For IBGP peers, the peer AS number must be the same as the autonomous system number of this router configured under the global level.</p> <p>This is a required command for each configured peer. This may be configured under the group level for all neighbors in a particular group.</p>
Default	No AS numbers are defined.
Parameters	<i>as-number</i> — The autonomous system number, expressed as a decimal integer.
	Values 1 — 65535

policy

Syntax	<pre>policy <i>policy-name</i> no policy</pre>
Context	config>service>vprn>bgp>next-hop-res
Description	<p>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.</p> <p>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.</p>
Default	no policy
Parameters	<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, etc.), the entire string must be enclosed within 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>router>bgp config>service>vprn>bgp
Description	<p>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.</p> <p>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.</p>
Default	no peer-tracking-policy
Parameters	<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, etc.), the entire string must be enclosed within 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	<p>This command configures the route preference for routes learned from the configured peer(s).</p> <p>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.</p> <p>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.</p> <p>The no form of the command used at the global level reverts to default value.</p> <p>The no form of the command used at the group level reverts to the value defined at the global level.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>
Default	170
Parameters	<i>preference</i> — The route preference, expressed as a decimal integer.
Values	1 — 255

prefix-limit

Syntax	prefix-limit <i>family limit</i> [log-only] [threshold percentage] [idle-timeout { <i>minutes</i> forever }] [post-import] no prefix-limit <i>family</i>
Context	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	<p>This command configures the maximum number of BGP routes that can be received from a peer before some 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.</p> <p>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.</p> <p>The no form of the command removes the prefix-limit.</p>
Default	No prefix limits for any address family.
Parameters	<p>log-only — Enables the 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.</p> <p><i>percent</i> — The threshold value (as a percentage) that triggers a warning message to be sent.</p> <p>Values 1 — 100</p> <p><i>family</i> — The address family to which the limit applies.</p> <p>Values ipv4 vpn-ipv4 ipv6 vpn-ipv6 mcast-ipv4 l2-vpn mvpn-ipv4 mdt-safi ms-pw flow-ipv4 route-target mcast-vpn-ipv4 mvpn-ipv6 flow-ipv6 evpn mcast-ipv6</p> <p><i>limit</i> — The number of routes that can be learned from a peer expressed as a decimal integer.</p> <p>Values 1 — 4294967295</p> <p><i>minutes</i> — Specifies duration in minutes before automatically re-establishing a session.</p> <p>Values 1 — 1024</p> <p>forever — Specifies that the session is reestablished only after clear router bgp command is executed.</p> <p>post-import — Specifies that the limit should be applied only to the number of routes that are accepted by import policies.</p>

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 OS software recognizes the set of AS numbers that are defined by IANA as private. These are AS numbers in the range 64512 through 65535, inclusive.
- The **no** form of 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.

leak-import

- Syntax** **leak-import** *policy-name* [*policy-name .. (up to 15 max)*]
no leak-import
- Context** config>service>vprn>bgp>rib-management>ipv4
config>service>vprn>bgp>rib-management>ipv6
- Description** This command specifies a chain of policy-statements (up to 15) that are evaluated against all the leakable IPv4 (or IPv6) BGP routes in all routing instances (VPRN or base). If a leak-import policy of routing instance T (the target routing instance) accepts a BGP route originally learned and marked as

BGP Commands

leakable by another routing instance S (the source routing instance), and the BGP route is valid in S (a best or non-best path) then a copy of the route is added to the BGP RIB-IN of instance T.

Default No default

Parameters *policy-name* — The name of a route policy statement (up to 32 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.

NOTE: 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	<p>This command designates the BGP peer as type internal or external.</p> <p>The type of internal indicates the peer is an IBGP peer while the type of external indicates that the peer is an EBGP peer.</p> <p>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.</p> <p>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.</p>
Default	no type — Type of neighbor is derived on the local AS specified.
Parameters	<p>internal — Configures the peer as internal.</p> <p>external — Configures the peer as external.</p>

updated-error-handling

Syntax	[no] updated-error-handling
Context	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
Description	<p>This command controls whether SROS utilizes the new neighbor-complete bit when processing optional transitive path attributes and advertising them to the associated BGP neighbor.</p> <p>This command also control if SROS utilizes the error handling mechanism for optional-transitive path attributes.</p>
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	<i>min-ttl-value</i> — Specify the minimum TTL value for an incoming BGP packet.

BGP Commands

Values 1 — 255

Default 1

ISIS Commands

isis

Syntax	[no] isis <i>isis-instance</i>
Context	config>service>vprn
Description	This command enables the ISIS protocol instance with the VPRN service. The no form of the command disables the ISIS protocol instance from the given VPRN service.
Default	no ISIS

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 <i>ieee-address</i> 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. You should shut/no shut the IS-IS instance to make the change operational.
Default	all-l1isis 01-80-C2-00-01-00
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. You should shut/no shut the IS-IS instance to make the change operational.
Default	all-l2isis 01-80-C2-00-02-11
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	<p>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.</p> <p>A maximum of 3 area addresses can be configured for the VPRN instance.</p> <p>NSAP addresses are divided into three parts. Only the area ID portion is configurable.</p> <ul style="list-style-type: none">• 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	confi>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** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

authentication-type

Syntax	authentication-type { password message-digest } no authentication
Context	config>service>vprn>isis config>service>vprn>isis>level
Description	<p>This command enables either simple password or message digest authentication or must go in either the global IS-IS or IS-IS level context.</p> <p>Both the authentication key and the authentication type on a segment must match. The authentication-key statement must also be included.</p> <p>Configure the authentication type on the global level in the config>router>isis context.</p> <p>Configure or override the global setting by configuring the authentication type in the config>router>isis>level context.</p> <p>The no form of the command disables authentication.</p>
Default	no authentication-type — No authentication type is configured and authentication is disabled.
Parameters	<p>password — Specifies that simple password (plain text) authentication is required.</p> <p>message-digest — Specifies that MD5 authentication in accordance with RFC2104 is required.</p>

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 <i>tag</i> 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	<p>This command configures export routing policies that determine the routes exported from the routing table to IS-IS.</p> <p>If no export policy is defined, non IS-IS routes are not exported from the routing table manager to IS-IS.</p> <p>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.</p> <p>If an aggregate command is also configured in the config>router context, then the aggregation is applied before the export policy is applied.</p> <p>Routing policies are created in the config>router>policy-options context.</p> <p>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.</p>
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> [log <i>percentage</i>] no export-limit
Context	config>service>vprn>isis
Description	<p>This command configures the maximum number of routes (prefixes) that can be exported into IS-IS from the route table for the VPRN instance.</p> <p>The no form of the command removes the parameters from the configuration.</p>
Default	no export-limit, the export limit for routes or prefixes is disabled..
Parameters	<p><i>number</i> — Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.</p> <p>Values 1 — 4294967295</p> <p>log percentage — Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.</p> <p>Values 1 — 100</p>

graceful-restart

Syntax	[no] graceful-restart
Context	config>service>vprn>isis
Description	<p>This command enables graceful-restart helper support for ISIS for the VPRN instance. The router will act as a helper to neighbors who are graceful-restart-capable and are restarting.</p> <p>When the control plane of a graceful-restart-capable router fails, the neighboring routers (graceful-restart helpers) temporarily preserve adjacency information so packets continue to be forwarded through the failed graceful-restart router using the last known routes. If the control plane of the graceful-restart router comes back up within the timer limits, then the routing protocols re-converge to minimize service interruption.</p> <p>The no form of the command disables graceful restart and removes all graceful restart configurations in the ISIS instance.</p>
Default	disabled

helper-disable

Syntax	[no] helper-disable
Context	config>service>vprn>isis>graceful-restart
Description	<p>This command disables the helper support for graceful restart for the VPRN instance.</p> <p>When graceful-restart is enabled, the router can be a helper (meaning that the router is helping a neighbor to restart) or be a restarting router or both. The router supports only helper mode. This facilitates the graceful restart of neighbors but will not act as a restarting router (meaning that the router will not help the neighbors to restart).</p> <p>The no helper-disable command enables helper support and is the default when graceful-restart is enabled.</p>
Default	disabled

hello-authentication

Syntax	[no] hello-authentication
Context	<pre>config>service>vprn>isis config>service>vprn>isis>interface config>service>vprn>isis>level</pre>
Description	<p>This command enables authentication of individual ISIS packets of HELLO type for the VPRN instance.</p> <p>The no form of the command suppresses authentication of HELLO packets.</p>

hello-padding

Syntax	[no] hello-padding {adaptive loose strict}
Context	config>service>vprn>isis
Description	<p>This command enables the padding of IS-IS hello messages for the VPRN instance, such that the message will be padded to a length of either 1492 or up to the maximum MTU value.</p> <p>The no form of the command disables IS-IS hello padding.</p>
Default	no hello-padding — hello padding is not configured
Parameters	<p>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.</p> <ol style="list-style-type: none"> 1. point-to-point interface—Hello PDUs are padded until the sender declares an adjacency on the link to be in 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 use the "classic" algorithm described in ISO 10589, this is when adjacency state is up. If the neighbor does not support the adjacency state TLV, then padding continues. 2. broadcast interface—Padding starts until at least one adjacency is up on the interface. <p>loose — Specifies the loose padding option; the loose padding may not be able to detect certain situations such as asymmetrical MTUs between the routing devices.</p> <ol style="list-style-type: none"> 1. 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. 2. broadcast interface—Padding starts until at least one adjacency (bcast only has up/down) is up on the interface. <p>strict — Specifies the strict padding option.</p> <ol style="list-style-type: none"> 1. 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. 2. 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	<p>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.</p> <p>This command enables ISIS to ignore the ATT bit and therefore suppress the installation of default routes.</p> <p>The no form of the command specifies that ISIS will not ignore LSP errors.</p>

iid-tlv-enable

Syntax	[no] iid-tlv-enable
Context	config>service>vprn>isis
Description	This command enables or disables the Instance Identifier (IID) TLV for the ISIS instance. When enabled, each I-IS instance marks its packets with the IID TLV containing its unique 16-bit IID for the routing domain. You should shut/no shut the isis instance to make the change operational.
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>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

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 — 65535

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	<p><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 (“ ”).</p> <p><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 (“ ”).</p> <p>This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.</p> <p>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.</p> <p>hash2 — Specifies the key is entered in a more complex encrypted form. If the hash2 parameter is not used, the less encrypted hash form is assumed.</p>

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	<p>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.</p> <p>To configure the hello authentication type at the interface context, use hello-authentication-type in the config>router>isis>interface context.</p> <p>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.</p> <p>The no form of the command disables hello authentication.</p>
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.
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.
Default	point-to-point — For IP interfaces on SONET channels. broadcast — For IP interfaces on Ethernet or unknown type physical interfaces.
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>if 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 <i>seconds</i> no hello-interval
Context	config>router>isis>if>level <i>level-number</i> config>service>vprn>isis>interface>level
Description	This command configures the interval in seconds between hello messages issued on this interface at this level. The no form of the command to reverts to the default value.
Default	3 — Hello interval default for the designated intersystem. 9 — Hello interval default for non-designated intersystems.
Parameters	<i>seconds</i> — The hello interval in seconds expressed as a decimal integer. Values 1 — 20000

hello-multiplier

Syntax	hello-multiplier <i>multiplier</i> no hello-multiplier
Context	config>router>isis>if>level <i>level-number</i> config>service>vprn>isis>interface>level
Description	This command configures the number of missing hello PDUs from a neighbor after the router declares the adjacency down. The no form of the command reverts to the default value.
Default	3 — The router can miss up to 3 hello messages before declaring the adjacency down.
Parameters	<i>multiplier</i> — The multiplier for the hello interval expressed as a decimal integer. Values 2 — 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 — 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 — 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	10 — A metric of 10 for the level on the interface is used.
Parameters	<i>metric</i> — The metric assigned for this level on this interface.
Values	1 — 16777215

passive

Syntax	[no] passive
Context	config>service>vprn>isis>interface config>service>vprn>isis>interface>level
Description	<p>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.</p> <p>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.</p> <p>The no form of the command removes the passive attribute.</p>
Special Cases	<p>Service Interfaces — Service interfaces (defined using the service-prefix command in config>router) are passive by default.</p> <p>All other Interfaces — All other interfaces are not passive by default.</p>
Default	<p>passive — Service interfaces are passive.</p> <p>no passive — All other interfaces are not passive.</p>

priority

Syntax	priority <i>number</i> no priority
Context	config>service>vprn>isis>interface>level
Description	<p>This command configures the priority of the IS-IS router interface for designated router election on a multi-access network.</p> <p>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.</p> <p>The no form of the command reverts to the default value.</p>
Default	64
Parameters	<i>number</i> — The priority for this interface at this level.
	Values 0 — 127

sd-offset

Syntax	sd-offset <i>offset-value</i> 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 — 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 — 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 will fail.

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**

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 will fail.

The **no** form of this command re-instates the default value for this command.

Default no loopfree-alternate-exclude

lsp-pacing-interval

Syntax **lsp-pacing-interval** *milliseconds*
no lsp-pacing-interval

Context config>service>vprn>isis>interface

Description This command configures the interval between LSP PDUs sent from this interface. To avoid bombarding adjacent neighbors with excessive data, pace the Link State Protocol Data Units (LSP's). If a value of zero is configured, no LSP's are sent from the interface.

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

Default **100** — LSPs are sent in 100 millisecond intervals.

ISIS Commands

Parameters *milliseconds* — The interval in milliseconds that IS-IS LSP's can be sent from the interface expressed as a decimal integer.

Values 0 — 65535

mesh-group

Syntax **mesh-group** {**value** | **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** — The 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 — 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 100

Parameters *seconds* — The interval in seconds that IS-IS LSPs can be sent on the interface.

1 — 65535

tag

Syntax	tag <i>tag</i> no tag
Context	config>service>vprn>isis>interface
Description	This command configures a route tag to the specified IP address of an interface.
Parameters	<i>tag</i> — [1 — 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.

ISIS Commands

The **no** form of the command disables support for IS-IS IPv6 TLVs for IPv6 routing.

Default disabled

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** *level-number*

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:

```
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
```

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.

Default level 1 or level 2

Parameters *level-number* — 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 — 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	10
Parameters	<i>metric</i> — Specifies the default metric for interfaces in the IPv4 multicast topology (MT4). 1 — 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	no default-ipv6-unicast-metric
Parameters	<i>ipv6-metric</i> — Specifies the default metric for IPv6 unicast. Values 1 — 16777215

default-metric

- Syntax** **default-metric** *ipv4 metric*
 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** 10

ipv4 metric — Specifies the default metric for IPv4 unicast.
- Values** 1 — 16777215

external-preference

- Syntax** **external-preference** *preference*
 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 the following table:

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 ^a
IS-IS Level 2 internal	18	Yes [*]

Route Type	Preference	Configurable
OSPF external	150	Yes
IS-IS Level 1 external	160	Yes
IS-IS Level 2 external	165	Yes
TMS	167	No
BGP	170	Yes
BGP	170	Yes

- a. 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 — 255

preference

Syntax **preference** *preference*
no preference

Context `cconfig>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 the following table:

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

Route Type	Preference	Configurable
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes ^a
IS-IS level 2 external	165	Yes ^a
BGP	170	Yes

a. 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 — 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 15 displays configuration combinations and the potential adjacencies that can be formed.

Table 15: 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.

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.

Default level-1/2

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** *link-group-name*

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 *link-group-name* — 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.
Parameters	<i>string</i> — Character string to be associated with the associated link-group.
Default	revert-members <i>oper-members</i>

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
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 — 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
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 *offset-value* — 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 — 6777215

ipv6-unicast-metric-offset

Syntax **ipv6-unicast-metric-offset** *offset-value*
no ipv6-unicast-metric-offset

Context config>service>vprn>isis>link-group

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 *offset-value* — 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 — 6777215

member

Syntax **[no] member** *interface-name*

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 *interface-name* — Name of the interface to be added or removed from the associated link-group.

oper-members

Syntax **oper-members [0-8]**
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

ISIS Commands

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 oper-members 0

revert-members

Syntax **revert-members [0-8]**
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.

loopfree-alternate

Syntax [**no**] **loopfree-alternate**

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-alternate

loopfree-alternate-exclude

Syntax **loopfree-alternate-exclude prefix-policy prefix-policy [prefix-policy... up to 5]**
no loopfree-alternate-exclude

Context config>service>vprn>isis

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 that prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.

The default action of the **loopfree-alternate-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 deletes the exclude prefix policy.

Parameters **prefix-policy** *prefix-policy* — Specifies the name of the prefix policy, up to 32 characters. The specified name must have been already defined.

lsp-lifetime

Syntax **lsp-lifetime** *seconds*
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 LSP’s 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 **1200** — LSPs originated by the router should be valid for 1200 seconds (20 minutes).

Parameters *seconds* — The time, in seconds, that the router wants the LSPs it originates to be considered valid by other routers in the domain.

Values 350 — 65535

lsp-mtu-size

Syntax **lsp-mtu-size** *size*
no lsp-mtu-size

Context config>service>vprn>isis

Description This command configures the LSP MTU *size*. If the *size* value is changed from the default using CLI or SNMP, then ISIS must be restarted in order 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 bounce IS-IS for the change to take effect.

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

Default	1492
Parameters	<i>size</i> — Specifies the LSP MTU size.
Values	490 — 9190

lsp-refresh-interval

Syntax	lsp-refresh-interval <i>seconds</i> no lsp-refresh-interval
Context	config>service>vprn>isis
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	600
Parameters	<i>seconds</i> — Specifies the refresh interval.
Values	150— 65535

lsp-wait

Syntax	lsp-wait <i>lsp-wait</i> [<i>lsp-initial-wait</i> [<i>lsp-second-wait</i>]]
Context	config>service>vprn>isis
Description	This command is used to customize the throttling of IS-IS LSP-generation. 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.
Parameters	<i>lsp-max-wait</i> — Specifies the maximum interval in seconds between two consecutive occurrences of an LSP being generated.
Values	1 — 120
Default	5
	<i>lsp-initial-wait</i> — Specifies the initial LSP generation delay in seconds.
Values	0 — 100
Default	0

lsp-second-wait — Specifies the hold time in seconds between the first and second LSP generation.

Values 1 — 100

Default 1

multi-topology

Syntax **[no] multi-topology**

Context config>service>vprn>isis

Description This command enables IS-IS multi-topology support.

Default disabled

ipv4-multicast

Syntax **[no] ipv4-multicast**

Context config>service>vprn>isis>multi-topology

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 **[no] ipv6-unicast**

Context config>service>vprn>isis>multi-topology

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] no overload
Context	config>service>vprn>isis
Description	<p>This command administratively sets the IS-IS router to operate in the overload state for a specific time period, in seconds, or indefinitely.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>The no form of the command causes the router to exit the overload state.</p>
Default	no overload
Parameters	<p><i>seconds</i> — The time, in seconds, that this router must operate in overload state.</p> <p>Default infinity (overload state maintained indefinitely)</p> <p>Values 60 — 1800</p>

overload-on-boot

Syntax	overload-on-boot [timeout seconds] no overload-on-boot
Context	config>service>vprn>isis
Description	<p>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:</p> <ol style="list-style-type: none"> 1. The timeout timer expires. 2. A manual override of the current overload state is entered with the config>router>isis>no overload command. <p>The no overload command does not affect the overload-on-boot function.</p> <p>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:</p> <p style="margin-left: 40px;">L1 LSDB Overload : Manual on boot (Indefinitely in overload)</p> <p style="margin-left: 40px;">L2 LSDB Overload : Manual on boot (Indefinitely in overload)</p> <p>This state can be cleared with the config>router>isis>no overload command.</p>

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.

Default no overload-on-boot

Use the **show router isis status** command to display the administrative and operational state as well as all timers.

Parameters **timeout seconds** — Configure the timeout timer for overload-on-boot in seconds.

Values 60 — 1800

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

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.

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> — The reference bandwidth in kilobits per second expressed as a decimal integer. Values 1 — 1000000000 <i>tbps Tera-bps</i> — The reference bandwidth in terabits per second expressed as a decimal integer. Values 1 — 4 <i>gbps Giga-bps</i> — The reference bandwidth in gigabits per second expressed as a decimal integer. Values 1 — 999 <i>mbps Mega-bps</i> — The reference bandwidth in megabits per second expressed as a decimal integer. Values 1 — 999 <i>kbps Kilo-bps</i> — reference bandwidth in kilobits per second expressed as a decimal integer. Values 1 — 999

rib-priority

Syntax	rib-priority { high } <i>prefix-list-name</i> tag <i>tag-value</i> 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	<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. tag <i>tag-value</i> — 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 – 4294967295

router-id

Syntax	router-id <i>ip-address</i> 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.

rsvp-shortcut

Syntax	[no] rsvp-shortcut
Context	config>service>vprn>isis
Description	This command enables the use of an RSVP-TE shortcut for resolving IGP routes by IS-IS or OSPF routing protocols. This command instructs IS-IS or OSPF to include RSVP LSPs originating on this node and terminating on the router-id of a remote node as direct links with a metric equal to the operational metric provided by MPLS. If the user enabled the relative-metric option for this LSP, IGP will apply the

shortest IGP cost between the endpoints of the LSP plus the value of the offset, instead of the LSP operational metric, when computing the cost of a prefix which is resolved to the LSP.

When a prefix is resolved to a tunnel next-hop, the packet is sent labeled with the label stack corresponding to the NHLFE of the RSVP LSP. Any network event causing an RSVP LSP to go down will trigger a full SPF computation which may result in installing a new route over another RSVP LSP shortcut as tunnel next-hop or over a regular IP next-hop.

When `rsvp-shortcut` is enabled at the IGP instance level, all RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in `configure>router>mpls>lsp>to`, corresponds to a router-id of a remote node. RSVP LSPs with a destination corresponding to an interface address or any other loopback interface address of a remote node are automatically not considered by IS-IS or OSPF. The user can, however, exclude a specific RSVP LSP from being used as a shortcut for resolving IGP routes by entering the `configure>router>mpls>lsp>no igp-shortcut` command.

The SPF in OSPF or IS-IS will only use RSVP LSPs as forwarding adjacencies, IGP shortcuts, or as endpoints for LDP-over-RSVP. These applications of RSVP LSPs are mutually exclusive at the IGP instance level. If the user enabled two or more options in the same IGP instance, then forwarding adjacency takes precedence over the shortcut application, which takes precedence over the LDP-over-RSVP application.

When ECMP is enabled on the system and multiple equal-cost paths exist for a prefix, the following selection criteria are used to pick up the set of next-hops to program in the data path:

- for a destination = tunnel-endpoint (including external prefixes with tunnel-endpoint as the next-hop):
 - Ⓜ select tunnel with lowest tunnel-index (ip next-hop is never used in this case)
- for a destination != tunnel-endpoint:
 - Ⓜ exclude LSPs with metric higher than underlying IGP cost between the endpoint of the LSP
 - Ⓜ prefer tunnel next-hop over ip next-hop
 - Ⓜ within tunnel next-hops:
 - i. select lowest endpoint to destination cost
 - ii. if same endpoint to destination cost, select lowest endpoint node router-id
 - iii. if same router-id, select lowest tunnel-index
 - Ⓜ within ip next-hops:
 - i. select lowest downstream router-id
 - ii. if same downstream router-id, select lowest interface-index
- Note though no ECMP is performed across both the IP and tunnel next-hops the tunnel endpoint lies in one of the shortest IGP paths for that prefix. In that case, the tunnel next-hop is always selected as long as the prefix cost using the tunnel is equal or lower than the IGP cost.

The ingress IOM will spray the packets for this prefix over the set of tunnel next-hops and IP next-hops based on the hashing routine currently supported for IPv4 packets.

This feature provides IGP with the capability to populate the multicast RTM with the prefix IP next-hop when both the `rsvp-shortcut` and the `multicast-import` options are enabled in IGP. The unicast RTM can still make use of the tunnel next-hop for the same prefix. This change is made possible with the enhancement by which SPF keeps track of both the direct first hop and the tunneled first hop of a node which is added to the Dijkstra tree.

The resolution and forwarding of IPv6 prefixes to IPv4 IGP shortcuts is not supported.
The **no** form of this command disables the resolution of IGP routes using RSVP shortcuts.

Default **no rsvp-shortcut**
Values

spf-wait

Syntax **[no] spf-wait** *spf-wait* [*spf-initial-wait* [*spf-second-wait*]]

Context config>service>vprn>isis

Description This command defines the maximum interval between two consecutive SPF calculations in seconds. 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 then next SPF will run after 4000 milliseconds, etc., until it reaches the *spf-wait* value. The SPF interval will stay at *spf-wait* value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to *spf-initial-wait*.

Default no spf-wait

Parameters *spf-wait* — Specifies the maximum interval in seconds between two consecutive spf calculations.

Values 1 — 120

Default 10

spf-initial-wait — Specifies the initial SPF calculation delay in milliseconds after a topology change.

Values 10 — 100000

Default 1000

spf-second-wait — Specifies the hold time in milliseconds between the first and second SPF calculation.

Values 1 — 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 dif-

ferent. 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</i> [<i>netmask</i>]} <i>level</i> [tag <i>tag</i>] no summary-address { <i>ip-prefix/mask</i> <i>ip-prefix</i> [<i>netmask</i>]}																				
Context	config>service>vprn>isis																				
Description	This command creates summary-addresses for the specified router or VPRN instance.																				
Default	none																				
Parameters	<i>ip-prefix/mask</i> — Specifies information for the specified IP prefix and mask length. <table border="0" style="margin-left: 20px;"> <tr> <td style="vertical-align: top;">Values</td> <td>ipv4-prefix: a.b.c.d (host bits must be 0)</td> </tr> <tr> <td></td> <td>ipv4-prefix-length: 0 — 32</td> </tr> <tr> <td></td> <td>ipv6-prefix: 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> <tr> <td></td> <td>ipv6-prefix-length: [0 — 128]</td> </tr> </table> <i>netmask</i> — The subnet mask in dotted decimal notation. <table border="0" style="margin-left: 20px;"> <tr> <td style="vertical-align: top;">Values</td> <td>0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)</td> </tr> </table> <i>level</i> — Specifies IS-IS level area attributes. <table border="0" style="margin-left: 20px;"> <tr> <td style="vertical-align: top;">Values</td> <td>level-1, level-2, level-1/2</td> </tr> </table> tag <i>tag</i> — Assigns a route tag to the summary address. <table border="0" style="margin-left: 20px;"> <tr> <td style="vertical-align: top;">Values</td> <td>1 — 4294967295</td> </tr> </table>	Values	ipv4-prefix: a.b.c.d (host bits must be 0)		ipv4-prefix-length: 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-length: [0 — 128]	Values	0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)	Values	level-1, level-2, level-1/2	Values	1 — 4294967295
Values	ipv4-prefix: a.b.c.d (host bits must be 0)																				
	ipv4-prefix-length: 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-length: [0 — 128]																				
Values	0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)																				
Values	level-1, level-2, level-1/2																				
Values	1 — 4294967295																				

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	suppress-attached-bit 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.

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	disabled

OSPF Commands

ospf

Syntax	[no] ospf
Context	config>service>vprn
Description	<p>This command enables access to the context to enable an OSPF protocol instance.</p> <p>When an OSPF instance is created, the protocol is enabled. To start or suspend execution of the OSPF protocol without affecting the configuration, use the no shutdown command.</p> <p>The no form of the command deletes the OSPF protocol instance removing all associated configuration parameters.</p>
Default	no ospf — The OSPF protocol is not enabled.

ospf3

Syntax	ospf3 [<i>instance-id</i>] [<i>router-id</i>] [no] ospf3 <i>instance-id</i>
Context	config>service>vprn
Description	<p>This command creates an OSPFv3 routing instance and then enters the associated context to configure associated protocol parameters.</p> <p>When an OSPFv3 instance is created, the protocol is enabled. To start or suspend execution of the OSPF.</p> <p>The no form of the command deletes the OSPFv3 protocol instance, removing all associated configuration parameters.</p>
Default	no default
Parameters	<p><i>instance-id</i> — Specify 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 64-95.</p> <p>Values 0 — 31: IPV6 unicast</p> <p>Values 64—95: IPV4 unicast</p>

advertise-router-capability

Syntax	advertise-router-capability { link area as } no advertise-router-capability
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	<p>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:</p> <ul style="list-style-type: none"> • 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 <p>The parameters (link, area & as) control the scope of the capabilities advertisements. The no form of this command, disables this capability.</p>
Default	no advertise-router-capability
Parameters	<p>link — Capabilities are only advertised over local link and not flooded beyond</p> <p>area — Capabilities are only advertised within the area of origin</p> <p>as — Capabilities are only advertised throughout the entire autonomous system</p>

area

Syntax	[no] area <i>area-id</i>				
Context	config>service>vprn>ospf config>service>vprn>ospf3				
Description	<p>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.</p> <p>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 etc., that are currently assigned to this area.</p>				
Default	no area — No OSPF areas are defined.				
Parameters	<p><i>area-id</i> — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.</p> <table> <tr> <td>Values</td> <td>0.0.0.0 — 255.255.255.255 (dotted decimal)</td> </tr> <tr> <td></td> <td>0 — 4294967295 (decimal integer)</td> </tr> </table>	Values	0.0.0.0 — 255.255.255.255 (dotted decimal)		0 — 4294967295 (decimal integer)
Values	0.0.0.0 — 255.255.255.255 (dotted decimal)				
	0 — 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> no area-range <i>ip-prefix/mask</i>
Context	config>service>vprn>ospf>area ospf>service>vprn>nssa config>service>vprn>ospf3>area
Description	<p>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.</p> <p>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.</p> <p>The no form of the command deletes the range (non) advertisement.</p>
Default	no area-range — No range of addresses are defined.
Special Cases	<p>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.</p> <p>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.</p>
Parameters	<p><i>ipv6-prefix/prefix-length</i> — The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.</p> <p>Values <i>ipv6-prefix</i> - x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d - x [0..FFFF]H - d [0..255]D <i>prefix-length</i> - [0..128]</p> <p><i>mask</i> — The subnet mask for the range expressed as a decimal integer mask length or in dotted decimal notation.</p> <p>Values 0 — 32 (mask length), 0.0.0.0 — 255.255.255.255 (dotted decimal)</p> <p>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.</p>

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

interface

- Syntax** **[no] interface** *ip-int-name* [**secondary**]
- 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** — No OSPF interfaces are defined.
- Parameters** *ip-int-name* — The IP interface name. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service vprn interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.
- If the IP interface name does not exist or does not have an IP address configured an error message will be returned.
- If the IP interface exists in a different area it will be moved to this area.
- secondary** — Allows multiple secondary adjacencies to be established over a single IP interface.

sham-link

- Syntax** **sham-link** *ip-int-name ip-address*
- Context** config>service>vprn>ospf>area
- 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-VPRN link and the backdoor.

OSPF Commands

- Parameters** *ip-int-name* — 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 within double quotes. If the IP interface name does not exist or does not have an IP address configured, an error message will be returned.
- ip-address* — 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.
- Note that 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.

authentication

- Syntax** **authentication bidirectional sa-name**
authentication inbound sa-name outbound sa-name
no authentication
- Context** config>service>vprn>ospf3>area>if
- Description** This command configures OPSFv3 confidentiality authentication.
- The no form of the command removes the SA name from the configuration.
- Parameters** **bidirectional sa-name** — Specifies the IPSec security association name in case the OSPFv3 traffic on the interface has to be authenticated.
- inbound sa-name** — Specifies the IPSec security association name in case the OSPFv3 traffic on the interface has to be authenticated.
- outbound sa-name** — 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	<p>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.</p> <p>Note that this command is not valid in the OSPF3 context.</p> <p>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.</p> <p>By default, no authentication key is configured.</p> <p>Note that this command is not supported in the OSPF context.</p> <p>The no form of the command removes the authentication key.</p>
Default	no authentication-key — No authentication key is defined.
Parameters	<p><i>authentication-key</i> — 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 (“ ”).</p> <p><i>hash-key</i> — 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 (“ ”).</p> <p>This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.</p> <p>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.</p> <p>hash2 — Specifies the key is entered in a more complex encrypted form. If the hash2 parameter is not used, the less encrypted hash form is assumed.</p>

authentication-type

Syntax	authentication-type { password message-digest } no authentication-type
Context	config>service>vprn>ospf>area>if config>service>vprn>ospf>area>virtual-link
Description	<p>This command enables authentication and specifies the type of authentication to be used on the OSPF interface, virtual-link, and sham-link.</p> <p>Note that this command is not valid in the OSPF3 context.</p>

OSPF Commands

Both simple **password** and **message-digest** authentication are supported.

By default, authentication is not enabled on an interface.

The **no** form of the command disables authentication on the interface.

Note that this command is not supported in the OSPF context.

Default	no authentication — 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>if config>service>vprn>ospf3>area>if
Description	
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	40
Special Cases	<p>OSPF Interface — If the dead-interval configured applies to an interface, then all nodes on the subnet must have the same dead interval.</p> <p>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.</p> <p>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.</p>
Parameters	<p><i>seconds</i> — The dead interval expressed as a decimal integer.</p> <p>Values 2 — 2147483647 seconds</p>

graceful-restart

Syntax	[no] graceful-restart
Context	config>service>vprn>ospf
Description	<p>This command enables or disables graceful-restart for VPRN OSPF.</p> <p>This command is not available for OSPF3.</p>

helper-disable

Syntax	helper-disable
Context	config>service>vprn>ospf>graceful-restart
Description	<p>This command disables the helper support for graceful restart.</p> <p>When graceful-restart is enabled, the router can be a helper (meaning that the router is helping a neighbor to restart) or be a restarting router or both. The helper-disable supports only helper mode. This facilitates the graceful restart of neighbors but will not act as a restarting router (meaning that the will not help the neighbors to restart).</p> <p>This command is not available for OSPF3.</p> <p>The no helper-disable command enables helper support and is the default when graceful-restart is enabled.</p>
Default	disabled

ignore-dn-bit

Syntax	[no] ignore-dn-bit
Context	config>service>vprn>ospf

OSPF Commands

```
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

hello-interval

- Syntax** **hello-interval** *seconds*
no hello-interval
- 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 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** *seconds* — The hello interval in seconds expressed as a decimal integer.
- Values** 1 — 65535

interface-type

- Syntax** **interface-type** {**broadcast** | **point-to-point**}
no interface-type
- Context** config>service>vprn>ospf>area>if
config>service>vprn>ospf3>area>if

Description	<p>This command configures the interface type to be either broadcast or point-to-point.</p> <p>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.</p> <p>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.</p> <p>The no form of the command reverts to the default value.</p>
Default	<p>point-to-point — If the physical interface is SONET.</p> <p>broadcast — If the physical interface is Ethernet or unknown.</p>
Special Cases	Virtual-Link — A virtual link is always regarded as a point-to-point interface and not configurable.
Parameters	<p>broadcast — Configures the interface to maintain this link as a broadcast network. To significantly improve adjacency forming and network convergence, a network should be configured as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.</p> <p>point-to-point — Configures the interface to maintain this link as a point-to-point link.</p>

loopfree-alternate-exclude

Syntax	[no] loopfree-alternate-exclude
Context	<pre>configure>service>vpn>ospf>area configure>service>vpn>ospf3>area configure>service>vpn>ospf>area>interface configure>service>vpn>ospf3>area>interface</pre>
Description	<p>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.</p> <p>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 will fail.</p> <p>The no form of this command re-instates the default value for this command.</p>
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	<p>This command enables filtering of outgoing OSPF LSAs on the selected OSPFv2 or OSPFv3 interface. Three filtering options are provided:</p> <ul style="list-style-type: none"> • 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. <p>The no form of this command disables OSPF LSA filtering (normal operation).</p>
Default	no lsa-filter-out

multicast-import

Syntax	[no] multicast-import
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	<p>This command enables the submission of routes into the multicast Route Table Manager (RTM) by OSPF.</p> <p>The no form of the command disables the submission of routes into the multicast RTM.</p>
Default	no multicast-import

message-digest-key

Syntax	message-digest-key <i>keyid</i> md5 [<i>key</i> <i>hash-key</i>] [hash] 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. Note that 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 — 255 md5 key — The MD5 key. The <i>key</i> can be any alphanumeric string up to 16 characters in length. md5 hash-key — 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 — The metric is based on reference-bandwidth setting and the link speed.
Parameters	metric — The metric to be applied to the interface expressed as a decimal integer. Values 1 — 65535

mtu

Syntax	mtu <i>bytes</i> 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.

Default **no mtu** — Uses the value derived from the MTU configured in the **config>port** context.

Parameters *bytes* — The MTU to be used by OSPF for this logical interface in bytes.

Values 512 — 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 Service interfaces defined in **config>router>service-prefix** are passive.

All other interfaces are not passive.

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 an election of the designated router on 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 Designated Router or Backup Designated Router.

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

Default **priority 1**

Parameters *number* — The interface priority expressed as a decimal integer. A value of 0 indicates the router is not eligible to be the Designated Router of Backup Designated Router on the interface subnet.

Values 0 — 255

retransmit-interval

Syntax **retransmit-interval** *seconds*
no retransmit-interval

Context config>service>vprn>ospf>area>if
config>service>vprn>ospf>area>virtual-link
config>service>vprn>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 acknowledgement has been received, the LSA will be retransmitted.

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

Default **retransmit-interval 5**

Parameters *seconds* — The retransmit interval in seconds expressed as a decimal integer.

Values 1 — 3600

rib-priority

Syntax	rib-priority { high } <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 — 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	10

Parameters *key-rollover-interval* — Specifies the time, in seconds, after which a key rollover will start.

Values 10 — 300

loopfree-alternate-exclude

Syntax **[no] loopfree-alternate-exclude**

Context config>service>vprn>ospf3>area

Description This command specifies whether or not the OSPF area should be excluded during LFA calculations. When enabled, the OSPF area is excluded from LFA calculations. When disabled (the default), the OSPF area is included in LFA calculations.

The **no** form of the command includes the OSPF area in LFA calculations.

Default disabled

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 originate-default-route

Context config>service>vprn>ospf>area>nssa
config>service>vprn>ospf3>area>nssa

OSPF Commands

Description	<p>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)</p> <p>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.</p> <p>The no form of the command disables origination of a default route.</p>
Default	no originate-default-route — A default route is not originated.
Parameters	<p>type-7 — Specifies a type 7 LSA should be used for the default route.</p> <p>Configure this parameter to inject a type-7 LSA default route instead the type 3 LSA into the NSSA configured with no summaries.</p> <p>To revert to a type 3 LSA, enter originate-default-route without the type-7 parameter.</p>
Default	Type 3 LSA for the default route.

redistribute-external

Syntax	[no] redistribute-external
Context	config>service>vprn>ospf>area>nssa config>service>vprn>ospf3>area>nssa
Description	<p>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.</p> <p>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.</p> <p>The no form of the command disables the default behavior to automatically redistribute external routes into the NSSA area from the NSSA ABR.</p>
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	<p>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.</p> <p>The no form of the command disables sending summary route advertisements and, for stub areas, only the default route is advertised by the ABR.</p>

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** *metric*
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 *metric* — The metric expressed as a decimal integer for the default route cost to be advertised into the stub area.

Values 1 — 16777215

virtual-link

Syntax **[no] virtual-link** *router-id* **transit-area** *area-id*

Context config>service>vprn>ospf>area
config>service>vprn>ospf3>area

Description This command configures a virtual link to connect 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 the picture below) then the area border routers (routers 1 and 2 in the picture below) must be connected via 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 the picture below). A virtual link can only be configured while in the area 0.0.0.0 context.

The *router-id* specified in this command must be associated with the virtual neighbor. The transit area cannot be a stub area or a Not So Stubby Area (NSSA).

The **no** form of the command deletes the virtual link.

Default No virtual link is defined.

Parameters *router-id* — The router ID of the virtual neighbor in IP address dotted decimal notation.

transit-area area-id — 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 30) 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 30: OSPF Areas

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 compliancy 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** *policy-name* [*policy-name...*]
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 polices 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** *policy-name* — The export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.
- The specified name(s) must already be defined.

external-db-overflow

- Syntax** **external-db-overflow** *limit interval*
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 *limit* value specifies the maximum number of non-default AS-external-LSA entries that can be stored in the link-state database (LSDB). Placing a limit on the non-default AS-external-LSAs in the LSDB protects the router from receiving an excessive number of external routes that consume excessive memory or CPU resources. If the number of routes reach or exceed the *limit*, the table is in an overflow state. When in an overflow state, the router 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	<p><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.</p> <p>Values -1 — 2147483647</p> <p><i>interval</i> — The number of seconds after entering an overflow state before attempting to process non-default AS-external-LSAs expressed as a decimal integer.</p> <p>Values 0 — 2147483647</p>

external-preference

Syntax	external-preference <i>preference</i> no external-preference
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	<p>This command configures the preference for OSPF external routes.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>The no form of the command reverts to the default value.</p>
Default	external-preference 150 — OSPF external routes have a default preference of 150.
Parameters	<i>preference</i> — 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 ^a
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

a. Preference for OSPF internal routes is configured with the **preference** command.

Values 1 — 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-alternate

Syntax	[no] loopfree-alternate
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.

OSPF Commands

The **no** form of this command disables the LFA computation by IGP SPF.

Default no loopfree-alternate

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.

Values 60 —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 <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no admin-group <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no admin-group
Context	config>router>interface>if-attribute config>service>ies>interface>if-attribute config>service>vprn>interface>if-attribute config>router>mpls>interface
Description	<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.</p> <p>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.</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>It should be noted that 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.</p> <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.</p>
Parameters	<i>group-name</i> — 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 <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no srlg-group <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no srlg-group
Context	config>router>interface>if-attribute config>service>ies>interface>if-attribute config>service>vprn>interface>if-attribute config>router>mpls>interface
Description	<p>This command configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface.</p> <p>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.</p> <p>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.</p>

It should be noted that 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.

lfa-policy-map

Syntax **lfa-policy-map route-nh-template** *template-name*
no lfa-policy-map

Context 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

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 will fail.

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 **loopfree-alternate-exclude prefix-policy** *prefix-policy* [*prefix-policy*... up to 5]
no loopfree-alternate-exclude

Context config>router>ospf
 config>router>ospf3
 config>router>isis
 config>service>vprn>ospf
 config>service>vprn>ospf3

- 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 that prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.
- The default action of the **loopfree-alternate-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 deletes the exclude prefix policy.
- Parameters** **prefix-policy** *prefix-policy* — Specifies the name of the prefix policy, up to 32 characters. The specified name must have been already defined.

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 <i>seconds</i>] no overload
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	<p>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:</p> <ul style="list-style-type: none"> • The timeout timer expires. • A manual override of the current overload state is entered with the no overload command. <p>The no overload command does not affect the overload-on-boot function.</p> <p>The no form of the command removes the overload-on-boot functionality from the configuration.</p>
Default	no overload-on-boot
Parameters	timeout <i>seconds</i> — Specifies the number of seconds to reset overloading.
	Values 60 —1800
	Default 60

preference

Syntax	preference <i>preference</i> no preference
Context	config>service>vprn>ospf config>service>vprn>ospf3
	<p>This command configures the preference for OSPF internal routes.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>The no form of the command reverts to the default value.</p>
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 ^a
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

a. Preference for OSPF internal routes is configured with the **preference** command.

Values 1 — 255

reference-bandwidth

Syntax	reference-bandwidth <i>reference-bandwidth</i> no reference-bandwidth
Context	config>service>vprn>ospf config>service>vprn>ospf3
Description	This command configures the reference bandwidth in kilobits per second (Kbps) that provides the reference for the default costing of interfaces based on their underlying link speed.

The default interface cost is calculated as follows:

$$\text{cost} = \text{reference-bandwidth} \div \text{bandwidth}$$

The default *reference-bandwidth* is 100,000,000 Kbps or 100 Gbps, so the default auto-cost metrics for various link speeds are as follows:

- 10 Mbs link default cost of 10000
- 100 Mbs link default cost of 1000
- 1 Gbps link default cost of 100
- 10 Gbps link default cost of 10

The **reference-bandwidth** command assigns a default cost to the interface based on the interface speed. To override this default cost on a particular interface, use the **metric** *metric* command in the **config>router>ospf>area>interface** *ip-int-name* context.

OSPF Commands

The **no** form of the command reverts the reference-bandwidth to the default value.

Default	reference-bandwidth 100000000 — Reference bandwidth of 100 Gbps.
Parameters	<i>reference-bandwidth</i> — The reference bandwidth in kilobits per second expressed as a decimal integer.
Values	1 — 1000000000

rib-priority

Syntax	rib-priority {high} <i>prefix-list-name</i> 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	<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.

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 enables 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 none

spf-wait

Syntax **spf-wait** *max-spf-wait* [*spf-initial-wait* [*spf-second-wait*]]
no spf-wait

Context config>service>vprn>ospf>timers
config>service>vprn>ospf3>timers

Description This command defines the maximum interval between two consecutive SPF calculations in milliseconds. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, then the next SPF will run after 2000 milliseconds, and then next SPF will run after 4000 milliseconds, etc., until it reaches the **spf-wait** value. The SPF interval will stay at the **spf-wait** value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to *spf-initial-wait*.

The timer must be entered in increments of 100 milliseconds. Values entered that do not match this requirement will be rejected.

Use the **no** form of this command to return to the default.

Default no spf-wait

Parameters *max-spf-wait* — Specifies the maximum interval in milliseconds between two consecutive SPF calculations.

Values 10 — 120000

Default 10000

spf-initial-wait — Specifies the initial SPF calculation delay in milliseconds after a topology change.

Values 10 — 100000

Default 1000

spf-second-wait — Specifies the hold time in milliseconds between the first and second SPF calculation.

Values 10 — 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

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 neighbors configured (lsa-generate) <i>lsa-second-wait</i> interval is equal or greater then the lsa-arrival timer configured here. Use the no form of this command to return to the default.
Default	no lsa-arrival
Parameters	<i>lsa-arrival-time</i> — Specifies the timer in milliseconds. Values entered that do not match this requirement will be rejected. Values 0 — 600000

lsa-generate

Syntax	lsa-generate <i>max-lsa-wait</i> [<i>lsa-initial-wait</i> [<i>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.

OSPF Commands

Default	no lsa-generate
Parameters	<i>max-lsa-wait</i> — Specifies the maximum interval, in milliseconds, between two consecutive occurrences of an LSA being generated. The timer must be entered as either 1 or in millisecond increments. Values entered that do not match this requirement will be rejected. Values 1 — 600000
Description	Default Values

VPRN Show Commands

egress-label

Syntax `egress-label egress-label1 [egress-label2]`

Context show>service

Description Display services using the range of egress labels.

If only the mandatory *egress-label1* parameter is specified, only services using the specified label are displayed.

If both *egress-label1* and *egress-label2* parameters are specified, the services using the range of labels X where *egress-label1* <= X <= *egress-label2* are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *egress-label1* — The starting egress label value for which to display services using the label range. If only *egress-label1* is specified, services only using *egress-label1* are displayed.

Values 0, 2049 — 131071

egress-label2 — The ending egress label value for which to display services using the label range.

Default The *egress-label1* value.

Values 2049 — 131071

Output **Show Service Egress Command Output** — The following table describes show service egress label output fields.

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.

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#
```

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 — 131071

end-label — The ending ingress label value for which to display services using the label range.

Default The *start-label* value.

Values 2048 — 131071

Output **Show Service Ingress-Label** — The following table describes show service ingress-label output fields:

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.

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#
```

sap-using

Syntax **sap-using** [**msap**] [**dyn-script**] [**description**]
sap-using [**sap** *sap-id*] [**vlan-translation** | **anti-spoof**] [**description**]
sap-using [**sap** *sap-id*]
sap-using **interface** [*ip-address* | *ip-int-name*]
sap-using [**ingress** | **egress**] **atm-td-profile** *td-profile-id*
sap-using [**ingress** | **egress**] **filter** *filter-id*
sap-using [**ingress** | **egress**] **qos-policy** *qos-policy-id*
sap-using **authentication-policy** *policy-name*

Context show>service

Description This command displays SAP information.
 If no optional parameters are specified, the command displays a summary of all defined SAPs.
 The optional parameters restrict output to only SAPs matching the specified properties.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common CLI Command Descriptions on page 849](#) for command syntax.
interface — Specifies matching SAPs with the specified IP interface.
ip-address — The IP address of the interface for which to display matching SAPs.
Values 1.0.0.0 — 223.255.255.255
ip-int-name — The IP interface name for which to display matching SAPs.
dyn-script — Displays dynamic service SAPs information.
ingress — Specifies matching an ingress policy.
egress — Specifies matching an egress policy.
qos-policy *qos-policy-id* — The ingress or egress QoS Policy ID for which to display matching SAPs.
Values 1 — 65535
atm-td-profile *td-profile-id* — Displays SAPs using this traffic description.
filter *filter-id* — The ingress or egress filter policy ID for which to display matching SAPs.
Values 1 — 65535
authentication-policy *policy name* — Specifies an existing authentication policy.

Output **Show Service SAP** — The following table describes show service SAP output fields:

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.

Label	Description (Continued)
E.QoS	The SAP egress QoS policy number specified on the egress SAP.
E.Mac/IP	The MAC or IP filter policy ID applied to the egress SAP
A.Pol	The accounting policy ID assigned to the SAP.
Adm	The desired state of the SAP.
Opr	The actual state of the SAP.

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
-----
/1/1:0/11 511111 2 none 2 none none Up Up
/1/1:0/12 511112 2 none 2 none none Up Up
/1/1:0/13 511113 2 none 2 none none Up Up
/1/1:0/14 511114 2 none 2 none none Up Up
/1/1:0/15 511115 2 none 2 none none Up Up
/1/1:0/16 511116 2 none 2 none none Up Up
/1/1:0/17 511117 2 none 2 none none Up Up
/1/1:0/18 511118 2 none 2 none none Up Up
/1/1:0/19 511119 2 none 2 none none Up Up
/1/1:0/20 511120 2 none 2 none none Up Up
/1/1:0/21 511121 2 none 2 none none Up Up
/1/1:0/22 511122 2 none 2 none none Up Up
/1/1:0/23 511123 2 none 2 none none Up Up
/1/1:0/24 511124 2 none 2 none none Up Up
/1/1:0/25 511125 2 none 2 none none Up Up ...
-----

*A:ALA-12#
```

sdp

Syntax `sdp [sdp-id | far-end ip-address] [detail | keep-alive-history]`

Context `show>service`

VPRN Show Commands

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

Parameters *sdp-id* — The SDP ID for which to display information.

Default All SDPs.

Values 1 — 17407

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

Default SDPs with any far-end IP address.

detail — Displays detailed SDP information.

Default SDP summary output.

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

Default SDP summary output.

Output **Show Service SDP** — The following table describes show service SDP output fields:

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.

Label	Description (Continued)
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.
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

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      180.0.0.2     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
```

VPRN Show Commands

```

=====
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                    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

```

Virtual Private Routed Network Services

```

Bw BookingFactor      : 100
Oper Max BW(Kbps)    : 0
Net-Domain           : default
Flags                : None
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
Default LSP           : Uknwn
EnforceDSTELspFc     : Disabled
Multicast LSP         : None
  
```

=====
 FC Mapping Table
 =====

```

FC Name              LSP Name
-----
  
```

```

No FC Mappings
  
```

 Segment Routing

```

ISIS                  : enabled
Oper Instance Id     : 0
LSP Id                : 524289
  
```

```

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
  
```

VPRN Show Commands

```

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#

```

sdp-using

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

Context `show>service`

Description Display services using SDP or far-end address options.

Parameters *sdp-id* — Displays only services bound to the specified SDP ID.

Values 1 — 17407

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

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

Default Services with any far-end IP address.

Output **Show Service SDP Using X** — The following table describes show service sdp-using output fields.

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.

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

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
```

service-using

Syntax `service-using [sdp sdp-id] [customer customer-id]`

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

- `ies` — Displays matching IES instances.
- `vprn` — Displays matching VPRN services.
- `mirror` — Displays mirror services.
- `sdp sdp-id` — Displays only services bound to the specified SDP ID.

Default Services bound to any SDP ID.

Values 1 — 17407

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

Default Services associated with an customer.

Values 1 — 2147483647

Output **Show Service Service-Using** — The following table describes show service service-using output fields:

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.

Sample Output

```
*A:ALA-12# show service service-using customer 10
=====
Services
=====
ServiceId   Type      Adm   Opr      CustomerId  Last Mgmt Change
-----
1           VPLS     Up    Up        10           09/05/2006 13:24:15
100        IES      Up    Up        10           09/05/2006 13:24:15
300        Epipe    Up    Up        10           09/05/2006 13:24:15
900        VPRN     Up    Up        2            11/04/2006 04:55:12
-----
Matching Services : 4
=====
*A:ALA-12#

*A:ALA-12# show service service-using epipe
=====
Services [epipe]
=====
ServiceId   Type      Adm   Opr      CustomerId  Last Mgmt Change
-----
6           Epipe    Up    Up        6            06/22/2006 23:05:58
7           Epipe    Up    Up        6            06/22/2006 23:05:58
8           Epipe    Up    Up        3            06/22/2006 23:05:58
103        Epipe    Up    Up        6            06/22/2006 23:05:58
-----
Matching Services : 4
=====
*A:ALA-12#

A:del14# show service service-using
```

```

=====
Services
=====
ServiceId   Type      Adm   Opr      CustomerId  Last Mgmt Change
-----
1           uVPLS     Up    Up        1            10/26/2006 15:44:57
2           Epipe     Up    Down      1            10/26/2006 15:44:57
10          mVPLS     Down  Down      1            10/26/2006 15:44:57
11          mVPLS     Down  Down      1            10/26/2006 15:44:57
100         mVPLS     Up    Up        1            10/26/2006 15:44:57
101         mVPLS     Up    Up        1            10/26/2006 15:44:57
102         mVPLS     Up    Up        1            10/26/2006 15:44:57
999         uVPLS     Down  Down      1            10/26/2006 16:14:33
-----
Matching Services : 8
-----
A:del14#

```

id

Syntax `id service-id {all | arp | base | fdb | labels | mfib | sap | sdp | split-horizon-group | stp}`

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.

all — Display detailed information about the service.

arp — Display ARP entries for the service.

base — Display basic service information.

fdb — Display FDB entries.

interface — Display service interfaces.

labels — Display labels being used by this service.

sap — Display SAPs associated to the service.

sdp — Display SDPs associated with the service.

split-horizon-group — Display split horizon group information.

stp — Display STP information.

all

Syntax `all`

Context show>service>id

Description Displays detailed information for all aspects of the service.

Show All Service-ID Output — The following table describes the show all service-id command output fields:

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.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.

<u>Split Horizon</u> Group	<u>Name of the split horizon group for this service</u>
Description	Description of the split horizon group.
Last Changed	The date and time of the most recent management-initiated change to this split horizon group.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.

Label	Description
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field. 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.

Label	Description
SAP Statistics	
Dropped	The number of packets or octets dropped.
Offered Hi Priority	The number of high priority packets, as determined by the SAP ingress QoS policy.
Offered Low Priority	The number of low priority packets, as determined by the SAP ingress QoS policy.
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
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.

Label	Description
Out profile for-warded	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Out profile dropped	The number of out-of-profile packets or octets discarded.
State	Specifies whether DHCP relay is enabled on this SAP.
Info Option	Specifies whether Option 82 processing is enabled on this SAP.
Action	Specifies the Option 82 processing on this SAP or interface: keep, replace or drop.
Circuit ID	Specifies whether the If index is inserted in circuit ID 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 signalling 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.

Sample Output

```

A:ALA-48# show service id 1 all
=====
Service Detailed Information
=====
Service Id       : 1                Vpn Id           : 0
Service Type     : VPRN
Customer Id      : 1
Last Status Change: 06/18/2007 10:07:01
Last Mgmt Change : 06/18/2007 10:07:01
Admin State      : Up                Oper State        : Up

Route Dist.      : 10001:1          VPRN Type        : regular
AS Number        : 10000            Router Id         : 10.10.10.103
ECMP              : Enabled          ECMP Max Routes   : 8
Max Routes       : 80                Auto Bind         : LDP
Vrf Target       : target:10001:1
Vrf Import       : vrfImpPolCust1
Vrf Export       : vrfExpPolCust1

SAP Count        : 2                SDP Bind Count    : 3
-----
Service Destination Points(SDPs)
-----
  Sdp Id 1:1 - (10.10.10.49)
-----
Description      : to-GRE-10.10.10.49
SDP Id           : 1:1                Type              : Spoke
VC Type          : n/a                VC Tag            : n/a
Admin Path MTU   : 0                  Oper Path MTU     : 0
Far End          : 10.10.10.49        Delivery           : GRE

Admin State      : Up                Oper State        : Down
Acct. Pol        : None               Collect Stats     : Disabled
Ingress Label    : n/a                Egress Label     : n/a
Ing mac Fltr     : n/a                Egr mac Fltr    : n/a
Ing ip Fltr      : n/a                Egr ip Fltr     : n/a
Admin ControlWord : Not Preferred      Oper ControlWord  : False
Admin BW(Kbps)   : 0                  Oper BW(Kbps)    : 0
Last Status Change: 06/18/2007 10:06:49
Last Mgmt Change : 06/18/2007 10:07:01
Signaling        : n/a
Class Fwding State : Down
Flags            : SdpOperDown
Peer Pw Bits     : None
Peer Fault Ip    : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None

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

Statistics        :
I. Fwd. Pkts.    : n/a                I. Dro. Pkts.    : n/a
I. Fwd. Octs.    : n/a                I. Dro. Octs.    : n/a
E. Fwd. Pkts.    : n/a                E. Fwd. Octets   : n/a

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-----

```

VP RN Show Commands

Sdp Id 1:10 -(10.10.10.49)

```

-----
Description      : to-GRE-10.10.10.49
SDP Id           : 1:10                               Type           : Spoke
VC Type         : n/a                               VC Tag         : n/a
Admin Path MTU  : 0                                 Oper Path MTU  : 0
Far End         : 10.10.10.49                       Delivery        : GRE

Admin State      : Up                               Oper State      : Down
Acct. Pol       : None                             Collect Stats   : Disabled
Ingress Label   : 0                               Egress Label   : 0
Ing mac Fltr    : n/a                             Egr mac Fltr   : n/a
Ing ip Fltr     : n/a                             Egr ip Fltr    : n/a
Ing ipv6 Fltr   : n/a                             Egr ipv6 Fltr  : n/a
Admin ControlWord : Not Preferred                 Oper ControlWord : False
Admin BW(Kbps)  : 0                               Oper BW(Kbps)   : 0
Last Status Change : 06/18/2007 10:06:49         Signaling      : n/a
Last Mgmt Change  : 06/18/2007 10:07:01
Class Fwding State : Down
Flags           : SdpOperDown
                  NoIngVCLabel NoEgrVCLabel
Peer Pw Bits    : None
Peer Fault Ip   : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None
  
```

KeepAlive Information :

```

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

Statistics :

```

I. Fwd. Pkts.    : 0                             I. Dro. Pkts.   : 0
I. Fwd. Octs.    : 0                             I. Dro. Octs.   : 0
E. Fwd. Pkts.    : 0                             E. Fwd. Octets  : 0
  
```

Associated LSP LIST :

SDP Delivery Mechanism is not MPLS

Sdp Id 3:4 -(10.10.10.105)

```

-----
SDP Id           : 3:4                               Type           : Spoke
VC Type         : n/a                               VC Tag         : n/a
Admin Path MTU  : 0                                 Oper Path MTU  : 0
Far End         : 10.10.10.105                     Delivery        : GRE

Admin State      : Up                               Oper State      : Down
Acct. Pol       : None                             Collect Stats   : Disabled
Ingress Label   : 3000                            Egress Label   : 2000
Ing mac Fltr    : n/a                             Egr mac Fltr   : n/a
Ing ip Fltr     : 10                               Egr ip Fltr    : 10
Ing ipv6 Fltr   : n/a                             Egr ipv6 Fltr  : n/a
Admin ControlWord : Not Preferred                 Oper ControlWord : False
Admin BW(Kbps)  : 0                               Oper BW(Kbps)   : 0
Last Status Change : 06/18/2007 10:06:49         Signaling      : n/a
Last Mgmt Change  : 06/18/2007 10:07:01
Class Fwding State : Down
Flags           : SdpOperDown
Peer Pw Bits    : None
Peer Fault Ip   : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None
  
```

```

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

```

```

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

```

```

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS

```

```

-----
Number of SDPs : 3
-----

```

```

Service Access Points
-----

```

```

SAP 1/1/21:0
-----

```

```

Service Id           : 1
SAP                  : 1/1/21:0           Encap                : q-tag
Dot1Q Ethertype      : 0x8100           QinQ Ethertype       : 0x8100

Admin State          : Up                 Oper State            : Down
Flags                : PortOperDown
Last Status Change  : 06/18/2007 10:06:49
Last Mgmt Change    : 06/18/2007 10:07:01
Admin MTU            : 1518              Oper MTU              : 1518
Ingress qos-policy   : 1                 Egress qos-policy    : 1
Shared Q plcy        : n/a              Multipoint shared     : Disabled
Ingr IP Fltr-Id     : n/a              Egr IP Fltr-Id       : n/a
Ingr Mac Fltr-Id    : n/a              Egr Mac Fltr-Id      : n/a
Ingr IPv6 Fltr-Id   : n/a              Egr IPv6 Fltr-Id     : n/a
tod-suite            : None              qinq-pbit-marking    : both
Egr Agg Rate Limit  : max

Multi Svc Site       : None
Acct. Pol            : None              Collect Stats         : Disabled

Anti Spoofing        : None              Nbr Static Hosts     : 0

```

```

-----
Sap Statistics
-----

```

```

Last Cleared Time    : N/A
                    Packets              Octets

Forwarding Engine Stats
Dropped              : 0                0
Off. HiPrio          : 0                0
Off. LowPrio         : 0                0
Off. Uncolor         : 0                0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio          : 0                0
Dro. LowPrio         : 0                0
For. InProf          : 0                0
For. OutProf         : 0                0

Queueing Stats(Egress QoS Policy 1)
Dro. InProf          : 0                0
Dro. OutProf         : 0                0
For. InProf          : 0                0
For. OutProf         : 0                0
-----

```

VPRN Show Commands

Sap per Queue stats

```
-----  
                          Packets          Octets  
Ingress Queue 1 (Unicast) (Priority)  
Off. HiPrio           : 0                0  
Off. LoPrio           : 0                0  
Dro. HiPrio           : 0                0  
Dro. LoPrio           : 0                0  
For. InProf           : 0                0  
For. OutProf          : 0                0
```

```
-----  
Egress Queue 1  
For. InProf           : 0                0  
For. OutProf          : 0                0  
Dro. InProf           : 0                0  
Dro. OutProf          : 0                0  
-----
```

SAP 1/2/4:0

```
-----  
Service Id           : 1  
SAP                  : 1/2/4:0          Encap           : q-tag  
Dot1Q Ethertype      : 0x8100          QinQ Ethertype  : 0x8100  
  
Admin State          : Up              Oper State       : Down  
Flags                : PortOperDown  
Last Status Change  : 06/18/2007 10:06:49  
Last Mgmt Change    : 06/18/2007 10:07:01  
Admin MTU            : 1518           Oper MTU         : 1518  
Ingress qos-policy   : 1             Egress qos-policy : 1  
Shared Q plcy        : n/a           Multipoint shared : Disabled  
Ingr IP Fltr-Id     : n/a           Egr IP Fltr-Id   : n/a  
Ingr Mac Fltr-Id    : n/a           Egr Mac Fltr-Id  : n/a  
Ingr IPv6 Fltr-Id   : n/a           Egr IPv6 Fltr-Id : n/a  
tod-suite            : None          qinq-pbit-marking : both  
Egr Agg Rate Limit  : max  
  
Multi Svc Site       : None  
Acct. Pol            : None          Collect Stats    : Disabled  
  
Anti Spoofing        : Ip-Mac         Nbr Static Hosts : 0  
-----
```

Subscriber Management

```
-----  
Admin State          : Down          MAC DA Hashing   : False  
Def Sub-Id           : None  
Def Sub-Profile      : None  
Def SLA-Profile      : None  
Def App-Profile      : None  
Sub-Ident-Policy     : None
```

```
Subscriber Limit     : 1  
Single-Sub-Parameters  
  Prof Traffic Only  : False  
  Non-Sub-Traffic    : N/A  
-----
```

Sap Statistics

```
-----  
Last Cleared Time    : N/A  
                          Packets          Octets  
Forwarding Engine Stats  
Dropped              : 0                0  
Off. HiPrio          : 0                0  
Off. LowPrio         : 0                0
```

Off. Uncolor : 0 0

Queueing Stats(Ingress QoS Policy 1)

Dro. HiPrio : 0 0
 Dro. LowPrio : 0 0
 For. InProf : 0 0
 For. OutProf : 0 0

Queueing Stats(Egress QoS Policy 1)

Dro. InProf : 0 0
 Dro. OutProf : 0 0
 For. InProf : 0 0
 For. OutProf : 0 0

 Sap per Queue stats

	Packets	Octets
Ingress Queue 1 (Unicast) (Priority)		
Off. HiPrio	: 0	0
Off. LoPrio	: 0	0
Dro. HiPrio	: 0	0
Dro. LoPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

Egress Queue 1

For. InProf	: 0	0
For. OutProf	: 0	0
Dro. InProf	: 0	0
Dro. OutProf	: 0	0

 Service Interfaces

Interface

If Name	: to-cel		
Admin State	: Up	Oper (v4/v6)	: Down/Down
Protocols	: None		
IP Addr/mask	: 11.1.0.1/24	Address Type	: Primary
IGP Inhibit	: Disabled	Broadcast Address	: Host-ones

 Details

If Index	: 2	Virt. If Index	: 2
Last Oper Chg	: 06/18/2007 10:07:01	Global If Index	: 96
SAP Id	: 1/1/21:0		
TOS Marking	: Trusted	If Type	: VPRN
SNTP B.Cast	: False		
MAC Address	: 14:30:01:01:00:15	Arp Timeout	: 14400
IP MTU	: 1500	ICMP Mask Reply	: True
Arp Populate	: Disabled	Host Conn Verify	: Enabled

Proxy ARP Details

Rem Proxy ARP	: Disabled	Local Proxy ARP	: Disabled
Policies	: none		

Proxy Neighbor Discovery Details

Local Pxy ND	: Disabled
Policies	: none

DHCP Details

Admin State	: Up	Lease Populate	: 1
Gi-Addr	: 11.1.0.1*	Gi-Addr as Src Ip	: Disabled

VPRN Show Commands

* = 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

Subscriber Authentication Details

Auth Policy : None

DHCP6 Relay Details

Admin State : Down Lease Populate : 0
Oper State : Down Nbr Resolution : Disabled
If-Id Option : None Remote Id : Disabled
Src Addr : Not configured

DHCP6 Server Details

Admin State : Down Max. Lease States : 8000

ICMP Details

Redirects : Number - 100 Time (seconds) - 10
Unreachables : Number - 100 Time (seconds) - 10
TTL Expired : Number - 100 Time (seconds) - 10

IPCP Address Extension Details

Peer IP Addr : Not configured
Peer Pri DNS Addr : Not configured
Peer Sec DNS Addr : Not configured

Interface

If Name : test
Admin State : Up Oper (v4/v6) : Down/Down
Protocols : IGMP PIM

IP Addr/mask : Not Assigned

Details

If Index : 3 Virt. If Index : 3
Last Oper Chg : 06/18/2007 10:07:01 Global If Index : 95
Port Id : n/a
TOS Marking : Trusted If Type : VPRN
SNTP B.Cast : False
MAC Address : Arp Timeout : 14400
IP MTU : 0 ICMP Mask Reply : True
Arp Populate : Disabled Host Conn Verify : Disabled

Proxy ARP Details

Rem Proxy ARP : Disabled Local Proxy ARP : Disabled
Policies : none

Proxy Neighbor Discovery Details

Local Pxy ND : Disabled
Policies : none

DHCP Details

Admin State : Down Lease Populate : 0
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

Admin State : Down Lease Populate : 0
 Oper State : Down Nbr Resolution : Disabled
 If-Id Option : None Remote Id : Disabled
 Src Addr : Not configured

DHCP6 Server Details

Admin State : Down Max. Lease States : 8000

ICMP Details

Redirects : Number - 100 Time (seconds) - 10
 Unreachables : Number - 100 Time (seconds) - 10
 TTL Expired : Number - 100 Time (seconds) - 10

IPCP Address Extension Details

Peer IP Addr : Not configured
 Peer Pri DNS Addr : Not configured
 Peer Sec DNS Addr : Not configured

 Interface

If Name : SpokeSDP
 Admin State : Up Oper (v4/v6) : Down/Down
 Protocols : None

IP Addr/mask : Not Assigned

Details

If Index : 4 Virt. If Index : 4
 Last Oper Chg : 06/18/2007 10:07:01 Global If Index : 94
 SDP Id : spoke-3:4
 TOS Marking : Trusted If Type : VPRN
 SNTP B.Cast : False
 MAC Address : 14:30:ff:00:00:00 Arp Timeout : 14400
 IP MTU : 0 ICMP Mask Reply : True
 Arp Populate : Disabled Host Conn Verify : Disabled

Proxy ARP Details

Rem Proxy ARP : Disabled Local Proxy ARP : Disabled
 Policies : none

Proxy Neighbor Discovery Details

Local Pxy ND : Disabled
 Policies : none

DHCP Details

Admin State : Down Lease Populate : 0
 Gi-Addr : Not configured Gi-Addr as Src Ip : Disabled
 Action : Keep Trusted : Disabled

DHCP Proxy Details

Admin State : Down
 Lease Time : N/A

VPRN Show Commands

Emul. Server : Not configured

Subscriber Authentication Details
Auth Policy : None

DHCP6 Relay Details

Admin State	: Down	Lease Populate	: 0
Oper State	: Down	Nbr Resolution	: Disabled
If-Id Option	: None	Remote Id	: Disabled
Src Addr	: Not configured		

DHCP6 Server Details

Admin State	: Down	Max. Lease States	: 8000
-------------	--------	-------------------	--------

ICMP Details

Redirects	: Number - 100	Time (seconds)	- 10
Unreachables	: Number - 100	Time (seconds)	- 10
TTL Expired	: Number - 100	Time (seconds)	- 10

IPCP Address Extension Details

Peer IP Addr : Not configured
Peer Pri DNS Addr : Not configured
Peer Sec DNS Addr : Not configured

Interface

If Name	: gizmo		
Admin State	: Up	Oper (v4/v6)	: Down/--
Protocols	: None		

IP Addr/mask : Not Assigned

Details

If Index	: 5	Virt. If Index	: 5
Last Oper Chg	: 06/18/2007 10:07:01	Global If Index	: 93
SDP Id	: spoke-1:10		
TOS Marking	: Trusted	If Type	: VPRN Red
Egress Filter	: none	Ingress Filter	: none
SNTP B.Cast	: False	QoS Policy	: 1
MAC Address	: 14:30:ff:00:00:00		
IP MTU	: 0	ICMP Mask Reply	: True

Interface

If Name	: test123		
Admin State	: Up	Oper (v4/v6)	: Down/--
Protocols	: None		

IP Addr/mask : Not Assigned

Details

If Index	: 6	Virt. If Index	: 6
Last Oper Chg	: 06/18/2007 10:07:01	Global If Index	: 92
Port Id	: n/a		
TOS Marking	: Trusted	If Type	: VPRN Red
Egress Filter	: none	Ingress Filter	: none
SNTP B.Cast	: False	QoS Policy	: 1
MAC Address	:		
IP MTU	: 0	ICMP Mask Reply	: True

Interface

```

-----
If Name           : test1
Admin State       : Up
Protocols         : None
Oper (v4/v6)     : Down/--

IP Addr/mask     : Not Assigned
-----
Details
-----
If Index          : 7
Last Oper Chg    : 06/18/2007 10:07:01
Port Id          : n/a
TOS Marking      : Trusted
Egress Filter    : none
SNTP B.Cast      : False
MAC Address      :
IP MTU           : 0
Virt. If Index   : 7
Global If Index  : 91
If Type          : VPRN Red
Ingress Filter   : none
QoS Policy       : 1
ICMP Mask Reply  : True
-----
Interface
-----
If Name           : bozocrown
Admin State       : Up
Protocols         : None
Oper (v4/v6)     : Down/--

IP Addr/mask     : Not Assigned
-----
Details
-----
If Index          : 8
Last Oper Chg    : 06/18/2007 10:07:01
Port Id          : n/a
TOS Marking      : Trusted
Egress Filter    : none
SNTP B.Cast      : False
MAC Address      :
IP MTU           : 0
Virt. If Index   : 8
Global If Index  : 90
If Type          : VPRN Red
Ingress Filter   : none
QoS Policy       : 1
ICMP Mask Reply  : True
-----
Interface
-----
If Name           : testabc
Admin State       : Up
Protocols         : None
Oper (v4/v6)     : Down/--

IP Addr/mask     : Not Assigned
-----
Details
-----
If Index          : 9
Last Oper Chg    : 06/18/2007 10:07:01
If Type          : VPRN Sub
Virt. If Index   : 9
Global If Index  : 89
-----
DHCP Details
Gi-Addr          : Not configured
Gi-Addr as Src Ip : Disabled
-----
Interface testabc group-interfaces
=====
Interface-Name   Adm      Opr (v4/v6)  Mode   Port/SapId
IP-Address                               PfxState
-----
bozo              Up        Down/--     VPRN G* n/a
-----
Group-Interfaces : 1
=====

```

VPRN Show Commands

* indicates that the corresponding row element may have been truncated.

Interface

```
-----  
If Name           : bozo  
Sub If Name       : testabc  
Red If Name       :  
Admin State       : Up           Oper (v4/v6)       : Down/--  
Protocols         : None  
-----
```

Details

```
-----  
If Index          : 10           Virt. If Index    : 10  
Last Oper Chg    : 06/18/2007 10:07:01 Global If Index   : 88  
Port Id          : n/a  
TOS Marking      : Trusted      If Type          : VPRN Grp  
SNTP B.Cast      : False  
MAC Address       :  
IP MTU           : 0           Arp Timeout      : 14400  
Arp Populate     : Disabled     ICMP Mask Reply   : True  
Host Conn Verify : Disabled
```

Proxy ARP Details

```
Rem Proxy ARP     : Disabled     Local Proxy ARP   : Enabled  
Policies          : none
```

Proxy Neighbor Discovery Details

```
Local Pxy ND      : Disabled  
Policies          : none
```

DHCP Details

```
Admin State       : Down         Lease Populate    : 1  
Gi-Addr          : Unknown       Gi-Addr as Src Ip : Disabled  
Action           : Keep         Trusted           : Disabled  
Match CircId     : Disabled
```

DHCP Proxy Details

```
Admin State       : Down  
Lease Time        : N/A  
Emul. Server      : Not configured
```

Subscriber Authentication Details

```
Auth Policy       : None
```

DHCP6 Relay Details

```
Admin State       : Down         Lease Populate    : 0  
Oper State        : Down         Nbr Resolution   : Disabled  
If-Id Option      : None        Remote Id         : Disabled  
Src Addr          : Not configured
```

DHCP6 Server Details

```
Admin State       : Down         Max. Lease States : 8000
```

ICMP Details

```
Redirects        : Number - 100   Time (seconds)   - 10  
Unreachables     : Number - 100   Time (seconds)   - 10  
TTL Expired      : Number - 100   Time (seconds)   - 10
```

IPCP Address Extension Details

```
Peer IP Addr      : Not configured  
Peer Pri DNS Addr : Not configured  
Peer Sec DNS Addr : Not configured
```

PPPoE Details

```

Last Mgmt Chg: 06/18/2007 10:06:49
Session limit      : 1                SAP session limit : 1
PPPoE Policy      : N/A
User DB           : N/A
=====
Service Access Point(Summary), Service 1 Interface bozo
=====
PortId            SvcId      Ing.  Ing.  Egr.  Egr.  Anti  Adm  Opr
                  QoS      Fltr  QoS  Fltr  Spoof
-----
No Service Access Point found.
=====
Interface
-----
If Name           : santa
Admin State       : Up                Oper (v4/v6)       : Down/--
Protocols         : None

IP Addr/mask      : Not Assigned
-----
Details
-----
If Index          : 11                Virt. If Index     : 11
Last Oper Chg    : 06/18/2007 10:07:01 Global If Index    : 87
If Type          : VPRN Sub

DHCP Details
Gi-Addr          : Not configured    Gi-Addr as Src Ip : Disabled
=====
Interface santa group-interfaces
=====
Interface-Name    Adm      Opr (v4/v6)  Mode      Port/SapId
  IP-Address                               PfxState
-----
interface         Up       Down/--      VPRN G*  1/2/4
-----
Group-Interfaces : 1
=====
* indicates that the corresponding row element may have been truncated.
-----
Interface
-----
If Name           : interface
Sub If Name       : santa
Red If Name       :
Admin State       : Up                Oper (v4/v6)       : Down/--
Protocols         : None

Details
-----
If Index          : 12                Virt. If Index     : 12
Last Oper Chg    : 06/18/2007 10:07:01 Global If Index    : 86
Group Port       : 1/2/4
TOS Marking      : Trusted           If Type            : VPRN Grp
SNTP B.Cast      : False
MAC Address      : 14:30:01:02:00:04 Arp Timeout        : 14400
IP MTU           : 1500             ICMP Mask Reply    : True
Arp Populate     : Disabled          Host Conn Verify   : Disabled

Proxy ARP Details
Rem Proxy ARP    : Disabled          Local Proxy ARP    : Enabled
Policies        : none

```

VPRN Show Commands

Proxy Neighbor Discovery Details

Local Pxy ND : Disabled
Policies : none

DHCP Details

Admin State : Down Lease Populate : 1
Gi-Addr : Unknown Gi-Addr as Src Ip : Disabled
Action : Keep Trusted : Disabled
Match CircId : Disabled

DHCP Proxy Details

Admin State : Down
Lease Time : N/A
Emul. Server : Not configured

Subscriber Authentication Details

Auth Policy : None

DHCP6 Relay Details

Admin State : Down Lease Populate : 0
Oper State : Down Nbr Resolution : Disabled
If-Id Option : None Remote Id : Disabled
Src Addr : Not configured

DHCP6 Server Details

Admin State : Down Max. Lease States : 8000

ICMP Details

Redirects : Number - 100 Time (seconds) - 10
Unreachables : Number - 100 Time (seconds) - 10
TTL Expired : Number - 100 Time (seconds) - 10

IPCP Address Extension Details

Peer IP Addr : Not configured
Peer Pri DNS Addr : Not configured
Peer Sec DNS Addr : Not configured

PPPoE Details

Last Mgmt Chg: 06/18/2007 10:06:49
Session limit : 1 SAP session limit : 1
PPPoE Policy : N/A
User DB : N/A

```
=====
Service Access Point(Summary), Service 1 Interface interface
=====
PortId          SvcId      Ing.  Ing.  Egr.  Egr.  Anti  Adm  Opr
                QoS      Fltr  QoS   Fltr  Spoof
-----
1/2/4:0         1          1    none  1     none  ip-mac Up  Down
=====
*#A:ALA-48#
```

The following output shows the IS-IS protocol.

*A:Dut-C# show service id 1 all

```
=====
Service Detailed Information
=====
Service Id      : 1                Vpn Id      : 0
Service Type    : VPRN
Name            : (Not Specified)
```


VPRN Show Commands

```
Last Mgmt Change   : 01/26/2015 11:22:14
Sub Type           : regular
Dot1Q Ethertype    : 0x8100
Split Horizon Group: (Not Specified)
QinQ Ethertype     : 0x8100

Admin MTU          : 1518
Ingr IP Fltr-Id   : 1
Ingr Mac Fltr-Id  : n/a
Ingr IPv6 Fltr-Id : n/a
BGP IPv4 FlowSpec : Disabled
BGP IPv6 FlowSpec : Disabled
tod-suite         : None

Q Frame-Based Acct : Disabled

Acct. Pol         : None

Anti Spoofing     : None
Avl Static Hosts  : 0
Calling-Station-Id : n/a

Application Profile: None
Transit Policy    : None
AARP Id          : None

Oper Group        : (none)
Host Lockout Plcy : n/a
Lag Link Map Prof : (none)

Oper MTU          : 1518
Egr IP Fltr-Id   : 1
Egr Mac Fltr-Id  : n/a
Egr IPv6 Fltr-Id : n/a

qinq-pbit-marking : both
Egr Agg Rate Limit: max
Limit Unused BW   : Disabled

Collect Stats     : Disabled

Dynamic Hosts     : Enabled
Tot Static Hosts  : 0

Monitor Oper Grp  : (none)
```

ETH-CFM SAP specifics

```
Tunnel Faults      : n/a
MC Prop-Hold-Timer : n/a
Squelch Levels     : 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)

Egress qos-policy  : 1
Egress Port QGrp   : (none)
Egr Port QGrp Inst: (none)
Multipoint shared  : Disabled
```

Sap Statistics

```
Last Cleared Time : N/A
```

```
                Packets      Octets
CPM Ingress     : 0           0
```

Forwarding Engine Stats

```
Dropped         : 0
Received Valid   : 0
Off. HiPrio     : 0
Off. LowPrio    : 0
Off. Uncolor    : 0
Off. Managed    : 0
```

Queueing Stats(Ingress QoS Policy 1)

Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

Queueing Stats(Egress QoS Policy 1)

Dro. InProf	: 0	0
Dro. OutProf	: 0	0
For. InProf	: 4	308
For. OutProf	: 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. InProf	: 4	308
For. OutProf	: 0	0
Dro. InProf	: 0	0
Dro. OutProf	: 0	0

Service Interfaces-----
Interface

If Name	: to_Ixial		
Admin State	: Up	Oper (v4/v6)	: Up/Up
Protocols	: None		
IP Addr/mask	: 1.3.9.3/24	Address Type	: Primary
IGP Inhibit	: Disabled	Broadcast Address	: Host-ones
HoldUp-Time	: 0	Track Srrp Inst	: 0
IPv6 Address	: 3ffe::103:903/120		
IPv6 Addr State	: PREFERRED		
CGA modifier	: (Not Specified)		
HoldUp-Time	: 0	Track Srrp Inst	: 0
Link Lcl Address	: fe80::200:ff:fe00:3/64		
Link Lcl State	: PREFERRED		
Description	: N/A		

Details

Description	: (Not Specified)		
If Index	: 2	Virt. If Index	: 2
Last Oper Chg	: 01/26/2015 11:22:14	Global If Index	: 262
Mon Oper Grp	: None		
Srrp En Rtng	: Disabled	Hold time	: N/A
SAP Id	: 1/1/4:1		
TOS Marking	: Trusted	If Type	: VPRN
SNTP B.Cast	: False		
MAC Address	: 00:00:00:00:00:03	Mac Accounting	: Disabled
Ingress stats	: Disabled	IPv6 DAD	: Enabled

VPRN Show Commands

```
TCP MSS V4      : 0
Arp Timeout    : 14400s
Arp Retry Timer : 5000ms
IP Oper MTU    : 1500
Arp Populate   : Disabled
Cflowd (unicast) : None
LdpSyncTimer   : None
LSR Load Balance : system
EGR Load Balance : both
TEID Load Balance : Disabled
SPI Load Balance : Disabled
uRPF Chk      : disabled
uRPF Ipv6 Chk : disabled
PTP HW Assist  : Disabled
Rx Pkts       : 0
Rx V4 Pkts    : N/A
Rx V6 Pkts    : N/A
Tx Pkts       : 2
Tx V4 Pkts    : 0
Tx V4 Discard Pkts: 0
Tx V6 Pkts    : 2
Tx V6 Discard Pkts: 0

TCP MSS V6      : 0
IPv6 Nbr ReachTime: 30s
ICMP Mask Reply : True
Host Conn Verify : Disabled
Cflowd (multicast): None

Rx Bytes       : 0
Rx V4 Bytes    : N/A
Rx V6 Bytes    : N/A
Tx Bytes       : 172
Tx V4 Bytes    : 0
Tx V4 Discard Byt*: 0
Tx V6 Bytes    : 172
Tx V6 Discard Byt*: 0

Proxy ARP Details
Rem Proxy ARP   : Disabled
Policies        : none

Local Proxy ARP : Disabled

Proxy Neighbor Discovery Details
Local Pxy ND    : Disabled
Policies        : none

Secure ND Details
Secure ND      : Disabled

DHCP no local server

DHCP Details
Description    : (Not Specified)
Admin State    : Down
Lease Populate : 0
Gi-Addr       : 1.3.9.3*
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

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 Time (seconds) - 10

Unreachables : Number - 100 Time (seconds) - 10

TTL Expired : Number - 100 Time (seconds) - 10

IPCP Address Extension Details

Peer IP Addr : Not configured

Peer Pri DNS Addr : Not configured

Peer Sec DNS Addr : Not configured

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:Dut-C#

The following output shows the OSPF protocol.

*A:Dut-C>config>service>vprn# show service id 1 all

=====
Service Detailed Information

=====
Service Id : 1 Vpn Id : 1
Service Type : VPRN
Name : XYZ Vprn 1
Description : (Not Specified)
Customer Id : 1 Creation Origin : manual
Last Status Change: 05/27/2015 17:58:34
Last Mgmt Change : 05/27/2015 17:58:35
Admin State : Up Oper State : Up

Route Dist. : 10.20.1.3:1 VPRN Type : regular
Oper Route Dist : 10.20.1.3:1
Oper RD Type : configured
AS Number : None Router Id : 10.20.1.3
ECMP : Enabled ECMP Max Routes : 1

Auto Bind Tunnel

Resolution : filter

Filter Protocol : sr-ospf

Max IPv6 Routes : No Limit

Ignore NH Metric : Disabled

Hash Label : Disabled

VPRN Show Commands

```

Vrf Target      : target:1:1
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 Inacti*: Disabled

SAP Count       : 1                SDP Bind Count    : 1

-----
ETH-CFM service specifics
-----
Tunnel Faults   : ignore

-----
VPRN service Network Specifics
-----
Ing Net QoS Policy : none
Ingress FP QGrp    : (none)                Ing FP QGrp Inst : (none)
-----

Service Destination Points(SDPs)
-----

Sdp Id 230:1  -(10.20.1.2)
-----
Description      : (Not Specified)
SDP Id           : 230:1                    Type              : Spoke
Spoke Descr      : (Not Specified)
VC Type          : n/a                      VC Tag            : n/a
Admin Path MTU   : 0                        Oper Path MTU     : 1582
Delivery         : MPLS
Far End          : 10.20.1.2
Tunnel Far End   : n/a                      LSP Types         : SR-OSPF
Hash Label       : Disabled                 Hash Lbl Sig Cap  : Disabled
Oper Hash Label  : Disabled

Admin State      : Up                       Oper State         : Up
Acct. Pol        : None                     Collect Stats     : Disabled
Ingress Label    : 262137                   Egress Label      : 262138
Ingr Mac Fltr-Id : n/a                      Egr Mac Fltr-Id  : n/a
Ingr IP Fltr-Id  : n/a                      Egr IP Fltr-Id   : n/a
Ingr IPv6 Fltr-Id : n/a                    Egr IPv6 Fltr-Id : n/a
BGP IPv4 FlowSpec : Disabled
BGP IPv6 FlowSpec : Disabled
Admin ControlWord : Not Preferred           Oper ControlWord  : False
BFD Template     : None
BFD-Enabled      : no                       BFD-Encap        : ipv4
Last Status Change : 05/27/2015 18:15:00      Signaling         : n/a
Last Mgmt Change  : 05/27/2015 17:58:35
Class Fwding State : Down
Flags            : None
Local Pw Bits    : None
Peer Pw Bits     : None
Peer Fault Ip    : None
Peer Vccv CV Bits : lspPing bfdFaultDet
Peer Vccv CC Bits : mplsRouterAlertLabel

```

Application Profile: None
 Transit Policy : None
 AARP Id : None

Ingress Qos Policy : (none)
 Ingress FP QGrp : (none)
 Ing FP QGrp Inst : (none)

Egress Qos Policy : (none)
 Egress Port QGrp : (none)
 Egr Port QGrp Inst: (none)

KeepAlive Information :

Admin State : Disabled
 Hello Time : 10
 Max Drop Count : 3

Oper State : Disabled
 Hello Msg Len : 0
 Hold Down Time : 10

Statistics :

I. Fwd. Pkts. : 0
 I. Fwd. Octs. : 0
 E. Fwd. Pkts. : 122

I. Dro. Pkts. : 0
 I. Dro. Octets : 0
 E. Fwd. Octets : 9372

 Control Channel Status

PW Status : disabled
 Peer Status Expire : false
 Request Timer : <none>
 Acknowledgement : false

Refresh Timer : <none>

 ETH-CFM SDP-Bind specifics

Squelch Levels : None

 RSVP/Static LSPs

Associated LSP List :
 No LSPs Associated

 Class-based forwarding :

Class forwarding : Disabled
 Default LSP : Uknwn

EnforceDSTELspFc : Disabled
 Multicast LSP : None

=====
 FC Mapping Table
 =====

FC Name LSP Name

 No FC Mappings

 Segment Routing

OSPF : enabled
 Oper Instance Id : 0

LSP Id : 524292

 Number of SDPs : 1

 Service Access Points

VPRN Show Commands

 SAP 1/1/4:1.1

```

Service Id          : 1
SAP                 : 1/1/4:1.1           Encap                : qinq
QinQ Dot1p         : Default
Description         : (Not Specified)
Admin State        : Up                   Oper State            : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 05/27/2015 17:49:33
Last Mgmt Change   : 05/27/2015 17:58:34
Sub Type           : regular
Dot1Q Ethertype    : 0x8100              QinQ Ethertype       : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU           : 1522                Oper MTU              : 1522
Ingr IP Fltr-Id    : n/a                 Egr IP Fltr-Id       : n/a
Ingr Mac Fltr-Id   : n/a                 Egr Mac Fltr-Id     : n/a
Ingr IPv6 Fltr-Id  : n/a                 Egr IPv6 Fltr-Id    : n/a
BGP IPv4 FlowSpec  : Disabled
BGP IPv6 FlowSpec  : Disabled
tod-suite          : None                 qinq-pbit-marking    : both
                                           Egr Agg Rate Limit  : max
                                           Limit Unused BW     : Disabled

Q Frame-Based Acct : Disabled

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)
  
```

 ETH-CFM SAP specifics

```

Tunnel Faults      : accept                AIS                    : Disabled
MC Prop-Hold-Timer : n/a
Squelch Levels     : None
  
```

 QOS

```

Ingress qos-policy : 2                     Egress qos-policy    : 2
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)
  
```

 Sap Statistics

```

Last Cleared Time   : N/A
  
```

	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 2)

Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

Queueing Stats(Egress QoS Policy 2)

Dro. InProf	: 0	0
Dro. OutProf	: 0	0
For. InProf	: 0	0
For. OutProf	: 129	11574

Sap per Queue stats

	Packets	Octets
--	---------	--------

Sap per Policer stats

	Packets	Octets
--	---------	--------

Ingress Policer 1 (Stats mode: minimal)

Off. All	: 0	0
Dro. All	: 0	0
For. All	: 0	0

Egress Policer 1 (Stats mode: minimal)

Off. All	: 129	11574
Dro. All	: 0	0
For. All	: 129	11574

Service Interfaces

Interface

If Name	: To_Ixia_1	Oper (v4/v6)	: Up/Down
Admin State	: Up		
Protocols	: OSPFv2		
IP Addr/mask	: 12.12.12.1/24	Address Type	: Primary
IGP Inhibit	: Disabled	Broadcast Address	: Host-ones
HoldUp-Time	: 0	Track Srrp Inst	: 0
Description	: N/A		

Details

Description	: (Not Specified)		
If Index	: 2	Virt. If Index	: 2
Last Oper Chg	: 05/27/2015 17:58:34	Global If Index	: 259
Mon Oper Grp	: None		
Srrp En Rtng	: Disabled	Hold time	: N/A

VPRN Show Commands

```
SAP Id : 1/1/4:1.1
TOS Marking : Trusted
SNTP B.Cast : False
MAC Address : 24:28:01:01:00:04
Ingress stats : Disabled
TCP MSS V4 : 0
ARP Timeout : 14400s
ARP Retry Timer : 5000ms
ARP Limit : Disabled
ARP Threshold : Disabled
ARP Limit Log Only: Disabled
IP MTU : (default)
IP Oper MTU : 1500
ARP Populate : Disabled
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
PTP HW Assist : Disabled

Rx Pkts : 0
Rx V4 Pkts : N/A
Rx V6 Pkts : N/A
Tx Pkts : 127
Tx V4 Pkts : 127
Tx V4 Discard Pkts: 0
Tx V6 Pkts : 0
Tx V6 Discard Pkts: 0
Mpls Rx Pkts : 0
Mpls Tx Pkts : 0

Proxy ARP Details
Rem 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
Gi-Addr : 12.12.12.1*
* = inferred gi-address from interface IP address

Action : Keep

DHCP Proxy Details
Admin State : Down
Lease Time : N/A
Emul. Server : Not configured

Subscriber Authentication Details
Auth Policy : None

DHCP6 Relay Details

If Type : VPRN
Mac Accounting : Disabled
IPv6 DAD : Enabled
TCP MSS V6 : 0
IPv6 Nbr ReachTime: 30s
IPv6 stale time : 14400s
IPv6 Nbr Limit : Disabled
IPv6 Nbr Threshold: Disabled
IPv6 Nbr Log Only : Disabled

ICMP Mask Reply : True
Host Conn Verify : Disabled
Cflowd (multicast): None

Rx Bytes : 0
Rx V4 Bytes : N/A
Rx V6 Bytes : N/A
Tx Bytes : 11430
Tx V4 Bytes : 11430
Tx V4 Discard Byt*: 0
Tx V6 Bytes : 0
Tx V6 Discard Byt*: 0
Mpls Rx Bytes : 0
Mpls Tx Bytes : 0

Local Proxy ARP : Disabled
Trusted : Disabled
```

```

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
Time (seconds) - 10
Time (seconds) - 10
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)
Ing FP QGrp Inst : (none)
Egress Port QGrp : (none)
Egr Port QGrp Inst: (none)

```

Interface

```

-----
If Name          : To_B_1
Admin State      : Up
Oper (v4/v6)     : Up/Down
Protocols        : OSPFv2
IP Addr/mask     : 11.11.11.1/24
Address Type     : Primary
IGP Inhibit      : Disabled
Broadcast Address : Host-ones
HoldUp-Time      : 0
Track Srrp Inst  : 0
Description      : N/A

```

Details

```

-----
Description      : (Not Specified)
If Index         : 3
Virt. If Index   : 3
Last Oper Chg   : 05/27/2015 18:15:00
Global If Index  : 260
Mon Oper Grp     : None
Srrp En Rtng    : Disabled
Hold time        : N/A
SDP Id           : spoke-230:1

```

Spoke-SDP Details

```

Admin State      : Up
Oper State       : Up

```

VPRN Show Commands

```
Hash Label          : Disabled          Hash Lbl Sig Cap   : Disabled
Oper Hash Label     : Disabled
Peer Fault Ip       : None
Local Pw Bits       : None
Peer Pw Bits        : None
Peer Vccv CV Bits   : lspPing bfdFaultDet
Peer Vccv CC Bits   : mplsRouterAlertLabel
Flags               : None

TOS Marking         : Trusted           If Type            : VPRN
SNTP B.Cast         : False
MAC Address         : 66:49: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
IP MTU             : 1500
IP Oper MTU        : 1500
ARP Populate       : Disabled          ICMP Mask Reply   : True
SHCV pol IPv4     : None              Host Conn Verify  : Disabled
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          : 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 Pkts          : 122                 Tx Bytes           : 9372
Tx V4 Pkts       : 0                  Tx V4 Bytes        : 0
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
Mpls Rx Pkts     : 0                  Mpls Rx Bytes      : 0
Mpls Tx Pkts     : 0                  Mpls Tx Bytes      : 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         : 11.11.11.1*       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

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

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.

VPLS SERVICE

*A:Dut-B# show service id 1 all

=====
 Service Detailed Information
 =====

Service Id	: 1	Vpn Id	: 0
Service Type	: VPLS		
Name	: vpls_1		
Description	: (Not Specified)		
Customer Id	: 1	Creation Origin	: manual

VPRN Show Commands

```
Last Status Change: 01/28/2015 21:59:54
Last Mgmt Change   : 01/28/2015 21:59:54
Etree Mode        : Disabled
Admin State       : Up
MTU               : 1514
SAP Count         : 1
Snd Flush on Fail : Disabled
Propagate MacFlush: Disabled
Allow IP Intf Bind: Disabled
Def. Gateway IP   : None
Def. Gateway MAC  : None
Temp Flood Time   : Disabled
Temp Flood Chg Cnt: 0
VSD Domain       : <none>
SPI load-balance  : Disabled

Oper State        : Up
Def. Mesh VC Id  : 1
SDP Bind Count   : 1
Host Conn Verify : Disabled
Per Svc Hashing  : Disabled

Temp Flood       : Inactive
```

BGP Information

Split Horizon Group specifics

ETH-CFM service specifics

```
Tunnel Faults      : ignore
V-Mep Extensions   : Enabled
```

Service Destination Points(SDPs)

```
Sdp Id 230:1  -(10.20.1.3)
```

```
Description      : (Not Specified)
SDP Id           : 230:1
Spoke Descr      : (Not Specified)
Split Horiz Grp  : (Not Specified)
Etree Root Leaf Tag: Disabled
VC Type          : Ether
Admin Path MTU   : 0
Delivery         : MPLS
Far End          : 10.20.1.3
Tunnel Far End   : n/a
Hash Label       : Disabled
Oper Hash Label  : Disabled

Type              : Spoke
Etree Leaf AC    : Disabled
VC Tag           : n/a
Oper Path MTU    : 1582
LSP Types        : SR-ISIS
Hash Lbl Sig Cap : Disabled

Admin State      : Up
Acct. Pol        : None
Ingress Label    : 262135
Ingr Mac Fltr-Id : n/a
Ingr IP Fltr-Id  : n/a
Ingr IPv6 Fltr-Id : n/a
Admin ControlWord : Not Preferred
BFD Template     : None
BFD-Enabled      : no
Last Status Change : 01/28/2015 22:00:07
Last Mgmt Change  : 01/28/2015 21:59:53
Endpoint         : N/A
PW Status Sig    : Enabled
Force Vlan-Vc    : Disabled

Oper State       : Up
Collect Stats    : Disabled
Egress Label     : 262135
Egr Mac Fltr-Id : n/a
Egr IP Fltr-Id  : n/a
Egr IPv6 Fltr-Id : n/a
Oper ControlWord : False
BFD-Encap       : ipv4
Signaling        : TLDP
Precedence       : 4
Force Qinq-Vc    : Disabled
```

```

Class Fwding State : Down
Flags               : None
Time to RetryReset : never
Mac Move           : Blockable
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing bfdFaultDet
Peer Vccv CC Bits  : mplsRouterAlertLabel
Retries Left       : 3
Blockable Level    : Tertiary

```

```

Application Profile: None
Transit Policy      : None
Max Nbr of MAC Addr: No Limit
Learned MAC Addr   : 0
OAM MAC Addr       : 0
Host MAC Addr      : 0
SPB MAC Addr       : 0
BGP EVPN Addr      : 0
Total MAC Addr     : 0
Static MAC Addr    : 0
DHCP MAC Addr      : 0
Intf MAC Addr      : 0
Cond MAC Addr      : 0
EVPN Static Addr   : 0

```

```

MAC Learning       : Enabled
MAC Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False
Oper Group         : (none)
Rest Prot Src Mac  : Disabled
Auto Learn Mac Prot: Disabled
Ing. Vlan Trans.   : 0
Discard Unkwn Srce: Disabled
Block On Mesh Fail: False
Monitor Oper Grp   : (none)
RestProtSrcMacAct : Disable

```

```

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

```

```

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

```

```

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

```

Control Channel Status

```

PW Status           : disabled
Peer Status Expire  : false
Request Timer       : <none>
Acknowledgement     : false
Refresh Timer       : <none>

```

```

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

```

ETH-CFM SDP-Bind specifics

```

V-MEP Filtering     : Disabled
Squelch Levels      : None

```

VPRN Show Commands

RSVP/Static LSPs

Associated LSP List :
No LSPs Associated

Class-based forwarding :

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

=====
FC Mapping Table
=====

FC Name LSP Name

No FC Mappings

Stp Service Destination Point specifics

Stp Admin State	: Up	Stp Oper State	: Down
Core Connectivity	: Down		
Port Role	: N/A	Port State	: Forwarding
Port Number	: 0	Port Priority	: 128
Port Path Cost	: 10	Auto Edge	: Enabled
Admin Edge	: Disabled	Oper Edge	: N/A
Link Type	: Pt-pt	BPDU Encap	: Dot1d
Root Guard	: Disabled	Active Protocol	: N/A
Last BPDU from	: N/A		
Designated 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
TC bit BPDUs rcvd	: 0	TC bit BPDUs tx	: 0
RST BPDUs rcvd	: 0	RST BPDUs tx	: 0

Number of SDPs : 1

Service Access Points

SAP 1/1/8:1.1

Service Id	: 1		
SAP	: 1/1/8:1.1	Encap	: qinq
QinQ Dot1p	: Default		
Description	: (Not Specified)		
Admin State	: Up	Oper State	: Up
Flags	: None		
Multi Svc Site	: None		
Last Status Change	: 01/28/2015 21:59:54		
Last Mgmt Change	: 01/28/2015 21:59:53		
Sub Type	: regular		
Dot1Q Ethertype	: 0x8100	QinQ Ethertype	: 0x8100
Split Horizon Group:	(Not Specified)		
Etree Root Leaf Tag:	Disabled	Etree Leaf Tag	: 0

```

Etree Leaf AC      : Disabled
Max Nbr of MAC Addr: No Limit
Learned MAC Addr   : 0
OAM MAC Addr       : 0
Host MAC Addr      : 0
SPB MAC Addr       : 0
BGP EVPN Addr      : 0
Admin MTU          : 1522
Ingr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id : n/a
tod-suite          : None

Q Frame-Based Acct : Disabled
ARP Reply Agent    : Disabled
Mac Learning       : Enabled
Mac Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None

Acct. Pol          : None

Anti Spoofing      : None
Avl Static Hosts   : 0
Calling-Station-Id : n/a

Application Profile: None
Transit Policy     : None

Oper Group         : (none)
Host Lockout Plcy : n/a
Lag Link Map Prof  : (none)
Cflowd            : Disabled
MCAC Policy Name   :
MCAC Max Unconst BW: no limit
MCAC In use Mand BW: 0
MCAC In use Opnl BW: 0
Use LAG port weight: no
Restr MacProt Src  : Disabled
Auto Learn Mac Prot: Disabled
Time to RetryReset : never
Mac Move          : Blockable
Egr MCast Grp     :
Auth Policy        : None

Total MAC Addr     : 0
Static MAC Addr    : 0
DHCP MAC Addr      : 0
Intf MAC Addr      : 0
Cond MAC Addr      : 0
EVPN Static Addr   : 0
Oper MTU           : 1522
Egr IP Fltr-Id    : n/a
Egr Mac Fltr-Id    : n/a
Egr IPv6 Fltr-Id  : n/a
qinq-pbit-marking  : both
Egr Agg Rate Limit: max
Limit Unused BW    : Disabled
Host Conn Verify   : Disabled
Discard Unkwn Srce: Disabled
Mac Pinning        : Disabled

Collect Stats      : Disabled

Dynamic Hosts      : Enabled
Tot Static Hosts   : 0

Monitor Oper Grp   : (none)

MCAC Const Adm St : Enable
MCAC Max Mand BW  : no limit
MCAC Avail Mand BW: unlimited
MCAC Avail Opnl BW: unlimited

Restr MacUnpr Dst  : Disabled
RestProtSrcMacAct : Disable
Retries Left       : 3
Blockable Level    : Tertiary

```

ETH-CFM SAP specifics

```

Tunnel Faults      : accept
MC Prop-Hold-Timer : n/a
Squelch Levels     : None

AIS                 : Disabled
V-MEP Filtering    : Disabled

```

Stp Service Access Point specifics

```

Stp Admin State    : Up
Core Connectivity   : Down
Port Role          : N/A
Port Number        : N/A
Port Path Cost     : 10
Admin Edge         : Disabled
Link Type          : Pt-pt

Stp Oper State     : Down
Port State         : Forwarding
Port Priority       : 128
Auto Edge          : Enabled
Oper Edge          : N/A
BPDU Encap        : Dot1d

```

VPRN Show Commands

```
Root Guard          : Disabled          Active Protocol    : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                Designated Port    : N/A

Forward transitions: 0                   Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                   Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                   TCN BPDUs tx      : 0
TC bit BPDUs rcvd  : 0                   TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                   RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                   MST BPDUs tx      : 0
-----

ARP host
-----

Admin State         : outOfService
Host Limit          : 1                   Min Auth Interval : 15 minutes
-----

QOS
-----

Ingress qos-policy : 2                   Egress qos-policy : 2
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)
-----

DHCP
-----

Description         : (Not Specified)
Admin State         : Down                Lease Populate     : 0
DHCP Snooping      : Down                Action             : Keep

Proxy Admin State   : Down
Proxy Lease Time    : N/A
Emul. Server Addr  : Not Configured
-----

Subscriber Management
-----

Admin State         : Down                MAC DA Hashing    : False
Def Sub-Id          : None
Def Sub-Profile     : None
Def SLA-Profile     : None
Def Inter-Dest-Id  : None
Def App-Profile     : None
Sub-Ident-Policy    : None

Subscriber Limit    : 1
Single-Sub-Parameters
  Prof Traffic Only : False
  Non-Sub-Traffic   : N/A

Static host management
MAC learn options   : N/A
-----

Sap Statistics
-----

Last Cleared Time   : N/A

                                Packets          Octets
CPM Ingress         : 0                      0
```

Forwarding Engine Stats

```

Dropped                : 0
Received Valid         : 0
Off. HiPrio            : 0
Off. LowPrio           : 0
Off. Uncolor           : 0
Off. Managed           : 0

```

Queueing Stats(Ingress QoS Policy 2)

```

Dro. HiPrio            : 0
Dro. LowPrio           : 0
For. InProf            : 0
For. OutProf           : 0

```

Queueing Stats(Egress QoS Policy 2)

```

Dro. InProf            : 0
Dro. OutProf           : 0
For. InProf            : 0
For. OutProf           : 0

```

Sap per Queue stats

```

-----
                                Packets                Octets
-----
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio                : 0                        0
Off. LowPrio                : 0                        0
Dro. HiPrio                 : 0                        0
Dro. LowPrio                : 0                        0
For. InProf                 : 0                        0
For. OutProf                : 0                        0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio                 : 0                        0
Off. LowPrio                : 0                        0
Off. Managed                : 0                        0
Dro. HiPrio                 : 0                        0
Dro. LowPrio                : 0                        0
For. InProf                 : 0                        0
For. OutProf                : 0                        0

Egress Queue 1
For. InProf                 : 0                        0
For. OutProf                : 0                        0
Dro. InProf                 : 0                        0
Dro. OutProf                : 0                        0

```

VPLS Spanning Tree Information

```

-----
VPLS oper state           : Up                      Core Connectivity : Down
Stp Admin State           : Down                    Stp Oper State    : Down
Mode                       : Rstp                      Vcp Active Prot.  : N/A

Bridge Id                  : 80:00:00:03:fa:32:16:62  Bridge Instance Id: 0
Bridge Priority             : 32768                          Tx Hold Count     : 6
Topology Change           : Inactive                       Bridge Hello Time : 2
Last Top. Change          : 0d 00:00:00                    Bridge Max Age    : 20
Top. Change Count         : 0                          Bridge Fwd Delay  : 15
MST region revision       : 0                          Bridge max hops   : 20
MST region name           :

```

VPRN Show Commands

```

Root Bridge      : N/A
Primary Bridge   : N/A

Root Path Cost   : 0
Rcvd Hello Time : 0
Root Priority     : 0

Root Forward Delay: 0
Root Max Age     : 0
Root Port       : N/A
  
```

Forwarding Database specifics

```

Service Id      : 1
Primary Factor   : 3
Mac Move Rate    : 2
Mac Move Retries : 3
Table Size      : 250
Learned Count    : 0
OAM MAC Count   : 0
Host MAC Count  : 0
Spb Count       : 0
BGP EVPN Count  : 0
Remote Age      : 900
High Watermark  : 95%
Mac Learning    : Enabled
Mac Aging       : Enabled
Mac Subnet Len  : 48

Mac Move        : Disabled
Secondary Factor : 2
Mac Move Timeout : 10

Total Count     : 0
Static Count    : 0
DHCP MAC Count  : 0
Intf MAC Count  : 0
Cond MAC Count  : 0
EVPN Static Cnt : 0
Local Age       : 300
Low Watermark   : 90%
Discard Unknown : Disabled
Relearn Only    : False
  
```

IGMP Snooping Base info

```

Admin State : Down
Querier     : No querier found
  
```

Sap/Sdp Id	Oper Stat	MRtr Port	Pim Port	Send Qrys	Max Grps	Max Srcs	Max Grp	MVR From-VPLS	Num Grps
------------	-----------	-----------	----------	-----------	----------	----------	---------	---------------	----------

sap:1/1/8:1.1	Up	No	No	No	None	None	None	Local	0
sdp:230:1	Up	No	No	No	None	None	None	N/A	0

MLD Snooping Base info

```

Admin State : Down
Querier     : No querier found
  
```

Sap/Sdp Id	Oper State	MRtr Port	Send Queries	Max Num Groups	MVR From-VPLS	Num Groups
------------	------------	-----------	--------------	----------------	---------------	------------

sap:1/1/8:1.1	Up	No	Disabled	No Limit	Local	0
sdp:230:1	Up	No	Disabled	No Limit	N/A	0

DHCP Summary, service 1

Sap/Sdp	Snoop	Used/Provided	Arp Reply Agent	Info Option	Admin State
---------	-------	---------------	-----------------	-------------	-------------

sap:1/1/8:1.1	No	0/0	No	Keep	Down
sdp:230:1	No	N/A	N/A	N/A	N/A

```

Number of Entries : 2
  
```

 ARP host Summary, service 1

Sap	Used	Provided	Admin State
sap:1/1/8:1.1	0	1	outOfService

 Number of SAPs : 1 0

=====
 WLAN Gateway specifics

Admin State : disabled
 Description : (Not Specified)
 SAP-template : (Not Specified)
 Last management change : (Not Specified)
 No associated WLAN Gateway interface VLAN tag ranges found.
 =====

=====
 Service VPLS Group Information

VPLS VXLAN, Ingress VXLAN Network Id: 0
 =====

Egress VTEP, VNI

VTEP Address	Egress VNI	Num. MACs	In Mcast List?	Oper State
No Matching Entries				

 Service Endpoints

No Endpoints found.

VPLS Sites

Site	Site-Id	Dest	Mesh-SDP	Admin	Oper	Fwdr
No Matching Entries						

=====
 *A:Dut-B#
 =====

EPIPE

*A:Dut-B# show service id 1 all

=====
 Service Detailed Information

Service Id	: 1	Vpn Id	: 0
Service Type	: Epipe		

VPRN Show Commands

```
Name : (Not Specified)
Description : (Not Specified)
Customer Id : 1 Creation Origin : manual
Last Status Change: 01/28/2015 22:05:35
Last Mgmt Change : 01/28/2015 22:05:22
Test Service : No
Admin State : Up Oper State : Up
MTU : 1514
Vc Switching : False
SAP Count : 1 SDP Bind Count : 1
Per Svc Hashing : Disabled
Force QTag Fwd : Disabled
```

BGP Information

ETH-CFM service specifics

Tunnel Faults : ignore

Service Destination Points(SDPs)

Sdp Id 230:1 -(10.20.1.3)

Description : (Not Specified)
SDP Id : 230:1 Type : Spoke
Spoke Descr : (Not Specified)
VC Type : Ether VC Tag : n/a
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
Hash Label : Disabled Hash Lbl Sig Cap : Disabled
Oper Hash Label : Disabled

Admin State : Up Oper State : Up
Acct. Pol : None Collect Stats : Disabled
Ingress Label : 262135 Egress Label : 262135
Ingr Mac Fltr-Id : n/a Egr Mac Fltr-Id : n/a
Ingr IP Fltr-Id : n/a Egr IP Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a Egr IPv6 Fltr-Id : n/a
Admin ControlWord : Not Preferred Oper ControlWord : False
Admin BW(Kbps) : 0 Oper BW(Kbps) : 0
BFD Template : None
BFD-Enabled : no BFD-Encap : ipv4
Last Status Change : 01/28/2015 22:05:35 Signaling : TLDP
Last Mgmt Change : 01/28/2015 22:05:22
Endpoint : N/A Precedence : 4
PW Status Sig : Enabled
Force Vlan-Vc : Disabled Force Qinq-Vc : Disabled
Class Fwding State : Down
Flags : None
Local Pw Bits : None
Peer Pw Bits : None
Peer Fault Ip : None
Peer Vccv CV Bits : lspPing bfdFaultDet
Peer Vccv CC Bits : mplsRouterAlertLabel

Application Profile: None
 Transit Policy : None
 Standby Sig Slave : False
 Block On Peer Fault: False
 Use SDP B-MAC : False

Ingress Qos Policy : (none)
 Ingress FP QGrp : (none)
 Ing FP QGrp Inst : (none)

Egress Qos Policy : (none)
 Egress Port QGrp : (none)
 Egr Port QGrp Inst: (none)

KeepAlive Information :

Admin State : Disabled
 Hello Time : 10
 Max Drop Count : 3

Oper State : Disabled
 Hello Msg Len : 0
 Hold Down Time : 10

Statistics :

I. Fwd. Pkts. : 0
 I. Fwd. Octs. : 0
 E. Fwd. Pkts. : 0

I. Dro. Pkts. : 0
 I. Dro. Octs. : 0
 E. Fwd. Octets : 0

 Control Channel Status

PW Status : disabled
 Peer Status Expire : false
 Request Timer : <none>
 Acknowledgement : false

Refresh Timer : <none>

 ETH-CFM SDP-Bind specifics

Squelch Levels : None

 RSVP/Static LSPs

Associated LSP List :
 No LSPs Associated

 Class-based forwarding :

Class forwarding : Disabled
 Default LSP : Uknwn

EnforceDSTELspFc : Disabled
 Multicast LSP : None

=====
 FC Mapping Table
 =====

FC Name LSP Name

 No FC Mappings

 Number of SDPs : 1

 Service Access Points

 SAP 1/1/8:1.1

Service Id : 1

VP RN Show Commands

```

SAP                : 1/1/8:1.1                Encap                : qinq
QinQ Dot1p        : Default
Description       : (Not Specified)
Admin State      : Up                        Oper State           : Up
Flags            : None
Multi Svc Site   : None
Last Status Change : 01/28/2015 22:05:22
Last Mgmt Change  : 01/28/2015 22:05:22
Sub Type         : regular
Dot1Q Ethertype  : 0x8100                    QinQ Ethertype      : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU         : 1522                      Oper MTU             : 1522
Ingr IP Fltr-Id  : n/a                       Egr IP Fltr-Id     : n/a
Ingr Mac Fltr-Id : n/a                       Egr Mac Fltr-Id    : n/a
Ingr IPv6 Fltr-Id : n/a                     Egr IPv6 Fltr-Id   : n/a
tod-suite        : None                       qinq-pbit-marking  : both
                                                        Egr Agg Rate Limit: max

Endpoint         : N/A
Q Frame-Based Acct : Disabled                Limit Unused BW     : Disabled
Vlan-translation : None

Acct. Pol        : None                       Collect Stats       : Disabled

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

-----
ETH-CFM SAP specifics
-----
Tunnel Faults    : accept                      AIS                  : Disabled
MC Prop-Hold-Timer : n/a
Squelch Levels   : None

-----
QOS
-----
Ingress qos-policy : 2                       Egress qos-policy : 2
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)

-----
Sap Statistics
-----
Last Cleared Time : N/A

                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 2)

```
Dro. HiPrio      : 0          0
Dro. LowPrio     : 0          0
For. InProf      : 0          0
For. OutProf     : 0          0
```

Queueing Stats(Egress QoS Policy 2)

```
Dro. InProf      : 0          0
Dro. OutProf     : 0          0
For. InProf      : 0          0
For. OutProf     : 0          0
```

Sap per Queue stats

	Packets	Octets
Ingress Queue 1 (Unicast) (Priority)		
Off. HiPrio	: 0	0
Off. LowPrio	: 0	0
Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0
Egress Queue 1		
For. InProf	: 0	0
For. OutProf	: 0	0
Dro. InProf	: 0	0
Dro. OutProf	: 0	0

Service Endpoints

No Endpoints found.

=====
VLL Sites
=====

Site	Site-Id	Dest	Admin	Oper	Fwdr
------	---------	------	-------	------	------

No Matching Entries

*A:Dut-B#

IES SPOKE

```
A:Dut-B#
*A:Dut-B#
*A:Dut-B#
*A:Dut-B# show service id 1 all
```

=====
Service Detailed Information
=====

```
Service Id      : 1          Vpn Id      : 0
Service Type    : IES
Name            : (Not Specified)
```

VPRN Show Commands

```
Description      : (Not Specified)
Customer Id      : 1
Creation Origin   : manual
Last Status Change: 01/28/2015 22:17:56
Last Mgmt Change : 01/28/2015 22:10:04
Admin State      : Up
Oper State       : Up
SAP Count        : 0
SDP Bind Count   : 1
```

ETH-CFM service specifics

```
Tunnel Faults    : ignore
```

Service Destination Points (SDPs)

```
Sdp Id 230:1 - (10.20.1.3)
```

```
Description      : (Not Specified)
SDP Id           : 230:1
Spoke Descr     : (Not Specified)
VC Type         : Ether
Admin Path MTU  : 0
Delivery        : MPLS
Far End         : 10.20.1.3
Tunnel Far End  : n/a
Hash Label      : Disabled
Oper Hash Label : Disabled

Type            : Spoke
VC Tag         : n/a
Oper Path MTU  : 1582
LSP Types      : SR-ISIS
Hash Lbl Sig Cap : Disabled

Admin State     : Up
Acct. Pol      : None
Ingress Label   : 262133
Ingr Mac Fltr-Id : n/a
Ingr IP Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a
BGP IPv4 FlowSpec : Disabled
BGP IPv6 FlowSpec : Disabled
Admin ControlWord : Not Preferred
BFD Template    : None
BFD-Enabled     : no
Last Status Change : 01/28/2015 22:17:56
Last Mgmt Change  : 01/28/2015 22:16:50
Class Fwding State : Down
Flags           : None
Local Pw Bits   : None
Peer Pw Bits    : None
Peer Fault Ip   : None
Peer Vccv CV Bits : lspPing bfdFaultDet
Peer Vccv CC Bits : mplsRouterAlertLabel

Oper State      : Up
Collect Stats   : Disabled
Egress Label    : 262133
Egr Mac Fltr-Id : n/a
Egr IP Fltr-Id : n/a
Egr IPv6 Fltr-Id : n/a
Oper ControlWord : False
BFD-Encap      : ipv4
Signaling       : TLDP
```

```
Application Profile: None
```

```
Transit Policy     : None
```

```
AARP Id           : None
```

```
Ingress Qos Policy : (none)
```

```
Ingress FP QGrp   : (none)
```

```
Ing FP QGrp Inst  : (none)
```

```
Egress Qos Policy : (none)
```

```
Egress Port QGrp : (none)
```

```
Egr Port QGrp Inst: (none)
```

```
KeepAlive Information :
```

```
Admin State       : Disabled
```

```
Hello Time        : 10
```

```
Max Drop Count    : 3
```

```
Oper State        : Disabled
```

```
Hello Msg Len     : 0
```

```
Hold Down Time    : 10
```

```

Statistics          :
I. Fwd. Pkts.      : 0          I. Dro. Pkts.      : 0
I. Fwd. Octs.      : 0          I. Dro. Octs.      : 0
E. Fwd. Pkts.      : 3          E. Fwd. Octets     : 180
    
```

Control Channel Status

```

PW Status           : disabled          Refresh Timer      : <none>
Peer Status Expire  : false
Request Timer       : <none>
Acknowledgement     : false
    
```

ETH-CFM SDP-Bind specifics

```

Squelch Levels     : 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

Number of SDPs : 1

Service Access Points

No Sap Associations

Service Interfaces

Interface

```

If Name            : iesSpokeToC
Admin State        : Up              Oper (v4/v6)       : Up/Down
Protocols          : None
IP Addr/mask       : 20.20.20.2/24    Address Type       : Primary
IGP Inhibit        : Disabled         Broadcast Address   : Host-ones
HoldUp-Time        : 0                Track Srrp Inst    : 0
Description        : N/A
    
```

Details

VPRN Show Commands

```

Description      : (Not Specified)
If Index         : 4          Virt. If Index    : 4
Last Oper Chg   : 01/28/2015 22:17:56 Global If Index : 257
Mon Oper Grp    : None
Srrp En Rtng   : Disabled    Hold time         : N/A
SDP Id          : spoke-230:1

Spoke-SDP Details
Admin State     : Up          Oper State        : Up
Hash Label      : Disabled    Hash Lbl Sig Cap  : Disabled
Oper Hash Label : Disabled
Peer Fault Ip   : None
Local Pw Bits   : None
Peer Pw Bits    : None
Peer Vccv CV Bits : lspPing bfdFaultDet
Peer Vccv CC Bits : mplsRouterAlertLabel
Flags           : None

TOS Marking     : Untrusted    If Type           : IES
SNTP B.Cast     : False        IES ID            : 1
MAC Address     : 00:03:fa:32:16:62 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
IP Oper MTU     : 1500         ICMP Mask Reply    : True
Arp Populate    : Disabled     Host Conn Verify   : Disabled
Cflowd (unicast) : None        Cflowd (multicast): None
LdpSyncTimer    : None
LSR Load Balance : system
EGR Load Balance : both
TEID Load Balance : Disabled
SPI Load Balance : Disabled
uRPF Chk        : disabled
uRPF Ipv6 Chk   : disabled
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 Pkts         : 3            Tx Bytes          : 180
Tx V4 Pkts      : 0            Tx V4 Bytes       : 0
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

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        : 20.20.20.2*   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

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

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:Dut-B#

*A:bkstim1618# show service id 2 all

 Service Detailed Information

Service Id : 2 Vpn Id : 0
 Service Type : VPRN

VPRN Show Commands

```
Name : (Not Specified)
Description : (Not Specified)
Customer Id : 1 Creation Origin : manual
Last Status Change: 08/21/2013 08:54:14
Last Mgmt Change : 08/21/2013 08:56:06
Admin State : Down Oper State : Down

Route Dist. : None VPRN Type : regular
AS Number : None Router Id : 18.18.18.18
ECMP : Enabled ECMP Max Routes : 1
Max IPv4 Routes : No Limit Auto Bind : None
Max IPv6 Routes : No Limit
Ignore NH Metric : Disabled
Hash 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

SAP Count : 0 SDP Bind Count : 0
-----
ETH-CFM service specifics
-----
Tunnel Faults : ignore
-----
Service Destination Points(SDPs)
-----
No Matching Entries
-----
Service Access Points
-----
No Sap Associations
-----
Service Interfaces
-----
No Interface Associations found.
-----
PTP Configuration
-----
Admin State : down Oper State : down
Peer Limit : 25
=====
```

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
tod-suite         : None
qinq-pbit-marking : both
Egr Agg Rate Limit: max

Endpoint           : N/A
Limit Unused BW   : Disabled

Vlan-translation  : None
Acct. Pol         : None
Collect Stats     : Disabled
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

```

```

-----
...

```

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 *policy name* — Specifies the subscriber authentication policy statistics to display.

sap sap-id — Specifies the SAP ID statistics to display. See [Common CLI Command Descriptions on page 849](#) for command syntax.

Sample Output

```

*A:ALA-1# show service id 11 authentication statistics
=====
Authentication statistics
=====
Interface / SAP           Authentication Successful  Authentication Failed
-----
-----

```

VPRN Show Commands

```

abc-11-90.1.0.254          1582          3
-----
Number of entries: 1
=====
*A:ALA-1#

```

arp

Syntax **arp** [*ip-address*] | [**mac** *ieee-address*] | [**sap** *sap-id*] | [**interface** *ip-int-name*] [**sdp** *sdp-id:vc-id*] [**summary**]

Context show>service>id

Description Displays the ARP table for the IES instance.

Parameters *ip-address* — Displays only ARP entries in the ARP table with the specified IP address.

Default All IP addresses.

mac *ieee-address* — Displays only ARP entries in the ARP table with the specified 48-bit MAC address. The MAC address can be expressed in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

Default All MAC addresses.

sap *sap-id* — Displays SAP information for the specified SAP ID. See [Common CLI Command Descriptions on page 849](#) for command syntax.

port id — Specifies matching service ARP entries associated with the specified IP interface.

ip-address — The IP address of the interface for which to display matching ARP entries.

Values 1.0.0.0 — 223.255.255.255

ip-int-name — The IP interface name for which to display matching ARPs.

Output **Show Service-ID ARP** — The following table describes show service-id ARP output fields.

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.

Sample Output

```
*A:ALA-12# show service id 2 arp
=====
ARP Table
=====
IP Address      MAC Address      Type    Age      Interface      Port
-----
190.11.1.1      00:03:fa:00:08:22 Other    00:00:00 ies-100-190.11.1 1/1/11:0
=====
*A:ALA-12#
```

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*]

Context show>service>id

Description This command displays ARP host related 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           Stdby
-----
128.128.1.2      00:80:00:00:00:01 2/1/5:2      00h04m41s
128.128.1.3      00:80:00:00:00:02 2/1/5:2      00h04m42s
128.128.1.4      00:80:00:00:00:03 2/1/5:2      00h04m43s
128.128.1.5      00:80:00:00:00:04 2/1/5:2      00h04m44s
128.128.1.6      00:80:00:00:00:05 2/1/5:2      00h04m45s
128.128.1.7      00:80:00:00:00:06 2/1/5:2      00h04m46s
128.128.1.8      00:80:00:00:00:07 2/1/5:2      00h04m47s
128.128.1.9      00:80:00:00:00:08 2/1/5:2      00h04m48s
128.128.1.10     00:80:00:00:00:09 2/1/5:2      00h04m49s
128.128.1.11     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 128.128.1.2 detail
=====
ARP hosts for service 2
=====
Service ID      : 2
IP Address      : 128.128.1.2
MAC Address     : 00:80:00:00:00:01
SAP             : 2/1/5:2
Remaining Time  : 00h04m58s

Sub-Ident       : "alu_1_2"
Sub-Profile-String : ""
```

VPRN Show Commands

```
SLA-Profile-String : ""
App-Profile-String : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
RADIUS-User-Name : "128.128.1.2"

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

Context show>service>id

Description Displays basic information about the service ID including service type, description, SAPs and SDPs.

Output **Show Service-ID Base** — The following table describes show service-id base output fields:

Label	Description
Service Id	The service identifier.
Vpn Id	Specifies the VPN ID assigned to the service.
Service Type	Specifies the type of service.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Adm	The desired state of the service.
Oper	The operating state of the service.
Mtu	The largest frame size (in octets) that the service can handle.
Def. Mesh VC Id	This object is only valid in services that accept mesh SDP bindings. It is used to validate the VC ID portion of each mesh SDP binding defined in the service.
SAP Count	The number of SAPs defined on the service.
SDP Bind Count	The number of SDPs bound to the service.
Identifier	Specifies the service access (SAP) and destination (SDP) points.
Type	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
AdmMTU	Specifies the desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented.
OprMTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented.
Opr	The operating state of the SDP.

Sample Output

```
*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)
```

VPRN Show Commands

```
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#
```

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.
ip-address — Specifies the ip-address of the neighbor.
port — Specifies the neighbor TCP port number use for this ANCP session.

Values 0 — 65535

association — Displays to what object the ANCP-string is associated.

statistics — Displays statistics information about an ANCP session known to the system.

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#

```

Note: 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
=====

```

VPRN Show Commands

```

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>gsmpr#

```

interface

Syntax `interface [ip-address | ip-int-name] [interface-type] [detail] [family]] summary]`

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 *ip-address* — The IP address of the interface for which to display information.

Values 1.0.0.0 — 223.255.255.255

ip-int-name — The IP interface name for which to display information.

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.

Output **Show Service-ID Interface** — The following table describes show service-id interface output fields:

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.

Label	Description (Continued)
Details	
If Index	The index corresponding to this interface. The primary index is 1. For example, all interfaces are defined in the Base virtual router context.
If Type	Specifies the interface type.
Port Id	Specifies the SAP'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.

Sample Output

```
*A:ALA-12# show service id 321 interface
=====
Interface Table
=====
Interface-Name          Type IP-Address      Adm   Opr   Type
-----
test                   Pri  190.11.1.1/24    Up    Up    IES
-----
Interfaces : 1
=====
*A:ALA-12#

A:ALA-49# show service id 88 interface detail
=====
Interface Table
=====
Interface
-----
If Name       : Sector A
Admin State   : Up
Oper State    : Down
Protocols     : None

IP Addr/mask  : Not Assigned
-----
Details
-----
Description :
```

VPRN Show Commands

```
If Index      : 26                               Virt. If Index : 26
SAP Id       : /1/1.2.2
TOS Marking  : Untrusted                         If Type       : IES
SNTP B.Cast  : False                             IES ID       : 88
MAC Address  : Not configured.                   Arp Timeout   : 14400
IP MTU       : 1500                              ICMP Mask Reply : True
Arp Populate : Disabled
Cflowd      : None
```

Proxy ARP Details

```
Proxy ARP    : Enabled                           Local Proxy ARP : Disabled
Policies     : ProxyARP
```

DHCP Details

```
Admin State : Up                               Lease Populate : 0
Action      : Keep                             Trusted       : Disabled
```

ICMP Details

```
Redirects   : Number - 100                     Time (seconds) - 10
Unreachables : Number - 100                   Time (seconds) - 10
TTL Expired  : Number - 100                   Time (seconds) - 10
```

Interface

```
-----
If Name     : test
Admin State : Up                               Oper State   : Down
Protocols   : None
IP Addr/mask : Not Assigned
-----
```

Details

```
-----
Description :
If Index    : 27                               Virt. If Index : 27
SAP Id     : /1/2:0
TOS Marking : Untrusted                         If Type       : IES
SNTP B.Cast : False                             IES ID       : 88
MAC Address : Not configured.                   Arp Timeout   : 14400
IP MTU     : 1500                              ICMP Mask Reply : True
Arp Populate : Disabled
Cflowd    : None
```

Proxy ARP Details

```
Proxy ARP    : Disabled                           Local Proxy ARP : Disabled
```

DHCP Details

```
Admin State : Up                               Lease Populate : 0
Action      : Keep                             Trusted       : Disabled
```

ICMP Details

```
Redirects   : Number - 100                     Time (seconds) - 10
Unreachables : Number - 100                   Time (seconds) - 10
TTL Expired  : Number - 100                   Time (seconds) - 10
```

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
               NoIngVCLabel NoEgrVCLabel

```

```

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                           Lease Populate  : 0
Oper State    : Down                           Nbr Resolution : Disabled
If-Id Option  : None                           Remote Id       : Disabled
Src Addr      : Not configured

```

DHCP6 Server Details

VPRN Show Commands

```
Admin State      : Down                               Max. Lease States: 8000

ICMP Details
Redirects       : Number - 100                       Time (seconds)    - 10
Unreachables   : Number - 100                       Time (seconds)    - 10
TTL Expired    : Number - 100                       Time (seconds)    - 10

IPCP Address Extension Details
Peer IP Addr*  : Not configured
Peer Pri DNS*  : Not configured
Peer Sec DNS*  : Not configured

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  -(2.2.2.2)
-----
Description   : (Not Specified)
SDP Id        : 1:555                               Type            : Spoke
Spoke Descr   : (Not Specified)
VC Type       : Ether                               VC Tag          : n/a
Admin Path MTU : 0                                 Oper Path MTU   : 1568
Far End       : 2.2.2.2                             Delivery        : MPLS
Tunnel Far End : n/a                               LSP Types       : RSVP
Hash Label    : Disabled                           Hash Lbl Sig Cap : Disabled
Oper Hash Label : Disabled

Admin State    : Up                               Oper State      : Up
Acct. Pol     : None                             Collect Stats   : Disabled
Ingress Label : 131065                            Egress Label   : 131059
Ingr Mac Fltr-Id : n/a                          Egr Mac Fltr-Id : n/a
Ingr IP Fltr-Id : n/a                          Egr IP Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a                        Egr IPv6 Fltr-Id : n/a
Admin ControlWord : Not Preferred                Oper ControlWord : False
Admin BW(Kbps)  : 0                               Oper BW(Kbps)   : 0
Last Status Change : 11/25/2010 13:06:14        Signaling       : TLDP
Last Mgmt Change  : 11/24/2010 13:00:48        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
```

```
.....
.....
=====
```

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.

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#
```

sap

Syntax `sap sap-id [detail]`

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 *sap-id* — The ID that displays SAPs for the service. See [Common CLI Command Descriptions on page 849](#) for command syntax.

detail — Displays detailed information for the SAP.

Output **Show Service-ID SAP** — The following table describes show service SAP fields:

Sample Output

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 Ether type 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, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEci-dOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapE-pipeRingNode.
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.

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.

```

*A:ALA-12# show service id 321 sap 1/1/4:0
=====
Service Access Points(SAP)
=====
Service Id       : 321
SAP              : 1/1/4:0
Dot1Q Ethertype : 0x8100
Admin State      : Up
Flags            : PortOperDown
                  SapIngressQoSMismatch
Last Status Change : 02/03/2007 12:58:37
Last Mgmt Change  : 02/03/2007 12:59:10
Admin MTU        : 1518
Ingress qos-policy : 100
Ingress Filter-Id : n/a
Multi Svc Site   : None
Acct. Pol        : None
Encap            : q-tag
QinQ Ethertype   : 0x8100
Oper State       : Down
Oper MTU         : 1518
Egress qos-policy : 1
Egress Filter-Id : n/a
Collect Stats    : Disabled
=====
*A:ALA-12#

```

*A:ALA-12# show service id 321 sap 1/1/4:0 detail

=====
 Service Access Points(SAP)
 =====

```

Service Id       : 321
SAP              : 1/1/4:0
Dot1Q Ethertype : 0x8100
Encap           : q-tag
QinQ Ethertype  : 0x8100

Admin State     : Up
Oper State      : Down
Flags           : PortOperDown
                  SapIngressQoSMismatch
Last Status Change : 02/03/2007 12:58:37
Last Mgmt Change  : 02/03/2007 12:59:10
Admin MTU       : 1518
Oper MTU        : 1518
Ingress qos-policy : 100
Egress qos-policy : 1
Ingress Filter-Id : n/a
Egress Filter-Id  : n/a
Multi Svc Site   : None
Acct. Pol        : None
Collect Stats    : Disabled
  
```

 Sap Statistics

	Packets	Octets
Forwarding Engine Stats		
Dropped	: 0	0
Off. HiPrio	: 0	0
Off. LowPrio	: 0	0
Off. Uncolor	: 0	0

Queueing Stats(Ingress QoS Policy 100)

Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

Queueing Stats(Egress QoS Policy 1)

Dro. InProf	: 0	0
Dro. OutProf	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

 Sap per Queue stats

	Packets	Octets
Ingress Queue 1 (Unicast) (Priority)		
Off. HiPrio	: 0	0
Off. LoPrio	: 0	0
Dro. HiPrio	: 0	0
Dro. LoPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

Ingress Queue 10 (Unicast) (Priority)

Off. HiPrio	: 0	0
Off. LoPrio	: 0	0
Dro. HiPrio	: 0	0
Dro. LoPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

'''

VPRN Show Commands

ATM SAP Configuration Information

Ingress TD Profile : 1 Egress TD Profile : 1
Alarm Cell Handling: Enabled AAL-5 Encap : VC-MUX

...

=====
*A:ALA-12#

sdp

Syntax **sdp** [*sdp-id* | **far-end** *ip-addr*] [**detail**]

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* — Displays only information for the specified SDP ID.

Default All SDPs.

Values 1 — 17407

far-end *ip-addr* — Displays only SDPs matching with the specified far-end IP address.

Default SDPs with any far-end IP address.

detail — Displays detailed SDP information.

Output **Show Service-ID SDP** — The following table describes show service-id SDP output fields:

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.

Label	Description (Continued)
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	he operational state of the keepalive process.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts.	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field. If the SDP type is GRE, then the following message displays: SDP delivery mechanism is not MPLS.

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
```

Virtual Private Routed Network Services

```

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

MAC Learning       : Enabled
MAC Aging          : Enabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled

MAC Learning       : Enabled
MAC Aging          : Enabled
L2PT Termination   : Disabled
MAC Pinning        : 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

I. Dro. Pkts.     : 0
I. Dro. Octs.     : 0
E. Fwd. Octets    : 0
MCAC Max Mand BW  : no limit
MCAC Avail Mand BW: unlimited
MCAC Avail Opnl BW: unlimited

Associated LSP LIST :
Lsp Name          : A_B_1
Admin State        : Up
Time Since Last Tr*: 00h26m35s

Oper State         : Up

Lsp Name          : A_B_2
Admin State        : Up
Time Since Last Tr*: 00h26m35s

Oper State         : Up

Lsp Name          : A_B_3
Admin State        : Up
Time Since Last Tr*: 00h26m34s

Oper State         : Up

Lsp Name          : A_B_4
Admin State        : Up
Time Since Last Tr*: 00h26m34s

Oper State         : Up

Lsp Name          : A_B_5
Admin State        : Up
Time Since Last Tr*: 00h26m34s

Oper State         : Up

Lsp Name          : A_B_6
Admin State        : Up
Time Since Last Tr*: 00h26m34s

Oper State         : Up

Lsp Name          : A_B_7
Admin State        : Up
Time Since Last Tr*: 00h26m34s

Oper State         : Up

Lsp Name          : A_B_8
Admin State        : Up
Time Since Last Tr*: 00h26m35s

Oper State         : Up

```

VPRN Show Commands

```

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#

```

sdp

Syntax

```

sdp sdp-id pw-port [pw-port-id]sdp
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>sdp

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* — Specifies the SDP ID for which to display information.

Default All SDPs.

Values 1 — 17407

pw-port-id — Specifies the pseudo-wire port identifier.

Values 1 — 10239

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

Default SDPs with any far-end IP address.

detail — Displays detailed SDP information.

Default SDP summary output.

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

Default SDP summary output.

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
```

VP RN Show Commands

```

=====
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
=====

```

pw-port

Syntax **pw-port** [*pw-port-id*] [*detail*]
pw-port sdp *sdp-id*
pw-port sdp none

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 — 10239

detail — Displays detailed port information that includes all the **pw-port** output fields.

sdp *sdp-id* — The SDP ID for which to display matching PW port information.

Values 1 — 17407

Output **Show PW-Port** — The following table describes **show pw-port** output fields:

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.

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      1        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
=====
```

filter-id

Syntax filter-id [*filter-id*]

Context show>service>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 *log-file-id* — Displays detailed information on the specified event file log.

Output **Log File Output** — The following table describes the output fields for a log file summary.

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.
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. none — 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. none — 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 [regex]] [ascending|descending]`

Context show>service>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.

Default Displays the event log summary

Values 1 — 99

severity severity-level — Displays only events with the specified and higher severity.

Default All severity levels

Values cleared, indeterminate, critical, major, minor, warning

application application — Displays only events generated by the specified application.

Default All applications

Values bgp, cflowd, chassis, debug, igmp, lldp, mirror, ospf, pim, port, snmp, system, user, vtr

expression — Specifies to use a regular expression as match criteria for the router instance string.

sequence 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**.

Default All sequence numbers

Values 1 — 4294967295

count count — Limits the number of log entries displayed to the *number* specified.

Default All log entries

Values 1 — 4294967295

router-instance — Specifies a router name up to 32 characters to be used in the display criteria.

subject 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.

regex — 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

Output **Show Log-ID Output** — The following table describes the log ID field output.

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.
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 's compact flash disks. 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.

Label	Description (Continued)
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>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 — 99

Output **SNMP Trap Group Output** — The following table describes SNMP trap group output fields.

Table 16: SNMP Trap Group Output Fields

Label	Description
Log-ID	The log destination ID for an event stream.
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 <code>snmpv1</code> , <code>snmpv2c</code> , <code>snmpv3</code> .
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 [<i>syslog-id</i>]
Context	show>service>log
Description	This command displays syslog event log destination summary information or detailed information on a specific syslog destination.
Parameters	<i>syslog-id</i> — Displays detailed information on the specified syslog event log destination.
	Values 1 — 10
Output	Syslog Event Log Destination Summary Output — The following table describes the syslog output fields.

Table 17: Show Log Syslog Output Fields

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 destination.

aggregate

Syntax	aggregate [active]
Context	show>router
Description	This command displays aggregated routes.
Parameters	active — This keyword filters out inactive aggregates.

Output Show Aggregate Output Fields — The following table describes router aggregate output fields.

Label	Description
Prefix	Displays the destination address of the aggregate route in dotted decimal notation.
Summary	Specifies whether the aggregate or more specific components are advertised.
AS Set	Displays an aggregate where the path advertised for the route consists of all elements contained in all paths that are being summarized.
Aggr AS	Displays the aggregator path attribute to the aggregate route.
Aggr IP-Address	The IP address of the aggregated route.
State	The operational state of the aggregated route.
No. of Aggregates	The total number of aggregated routes.

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          Aggr IP-Address  Aggr AS
                Summary          AS Set           State
                NextHop          NextHopType
-----
1.2.3.0/24      False           0.0.0.0          0
                2.2.2.2
                Indirect

2.2.0.0/16      False           0.0.0.0          0
                None
                Active

-----
No. of Aggregates: 2
=====

*A:CPM133>config>router# show router aggregate
=====
Aggregates (Router: Base)
```

VPRN Show Commands

```

=====
Prefix                               Aggr IP-Address  Aggr AS
  Summary                             AS Set          State
  NextHop                             Community       NextHopType
-----
10.0.0.0/8                            0.0.0.0         0
  False                               False           Inactive
                                       100:33         Blackhole
-----
No. of Aggregates: 1
=====

```

arp

Syntax `arp [ip-address | ip-int-name | mac ieee-mac-addr]`

Context `show>router`

Description This command displays the router ARP table sorted by IP address.

If no command line options are specified, all ARP entries are displayed.

Parameters *ip-addr* — Only displays ARP entries associated with the specified IP address.

ip-int-name — Only displays ARP entries associated with the specified IP interface name.

mac ieee-mac-addr — Only displays ARP entries associated with the specified MAC address.

Output **ARP Table Output** — The following table describes ARP table output fields:

Label	Description
IP Address	The IP address of the ARP entry.
MAC Address	The MAC address of the ARP entry.
Expiry	The age of the ARP entry.
Type	Dyn — The ARP entry is a dynamic ARP entry. Inv — The ARP entry is an inactive static ARP entry (invalid). Oth — The ARP entry is a local or system ARP entry. Sta — The ARP entry is an active static ARP entry.
Interface	The IP interface name associated with the ARP entry.
No. of ARP Entries	The number of ARP entries displayed in the list.

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#

```

damping

Syntax **damping** [*ip-prefix/mask* | *ip-address*] [**detail**]
damping [*damp-type*] [**detail**]

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/mask* — Displays damping information for the specified IP prefix and mask length.

ip-address — Displays damping entry for the best match route for the specified IP address.

damp-type — Displays damping type for the specified IP address.

decayed — Displays damping entries that are decayed but are not suppressed.

history — Displays damping entries that are withdrawn but have history.

suppressed — Displays damping entries suppressed because of route damping.

detail — Displays detailed information.

Output **Show Damping Output Fields** — The following table describes BGP damping output fields:

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured or inherited local AS for the specified peer group. If not configured, then it is the same value as the AS.
Network	Route IP prefix and mask length for the route.
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.

Label	Description (Continued)
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.

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 12.149.7.0/24       10.0.28.1     00h00m00s     60203 65001 19855 3356
                                     1239 22406
si   24.155.6.0/23       10.0.28.1     00h43m41s     60203 65001 19855 3356
                                     2914 7459
si   24.155.8.0/22       10.0.28.1     00h38m31s     60203 65001 19855 3356
                                     2914 7459
si   24.155.12.0/22      10.0.28.1     00h35m41s     60203 65001 19855 3356
                                     2914 7459
si   24.155.22.0/23      10.0.28.1     00h35m41s     60203 65001 19855 3356
                                     2914 7459
si   24.155.24.0/22      10.0.28.1     00h35m41s     60203 65001 19855 3356
                                     2914 7459
si   24.155.28.0/22      10.0.28.1     00h34m31s     60203 65001 19855 3356
                                     2914 7459
si   24.155.40.0/21      10.0.28.1     00h28m24s     60203 65001 19855 3356
                                     7911 7459
si   24.155.48.0/20      10.0.28.1     00h28m24s     60203 65001 19855 3356
                                     7911 7459
ud*i 61.8.140.0/24     10.0.28.1     00h00m00s     60203 65001 19855 3356
```

VPRN Show Commands

```

                                4637 17447
ud*i 61.8.141.0/24      10.0.28.1      00h00m00s 60203 65001 19855 3356
                                4637 17447
ud*i 61.9.0.0/18       10.0.28.1      00h00m00s 60203 65001 19855 3356
                                3561 9658 6163
. . .
ud*i 62.213.184.0/23   10.0.28.1      00h00m00s 60203 65001 19855 3356
                                6774 6774 9154

```

 *A: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 : 12.149.7.0/24

```

-----
Network      : 12.149.7.0/24      Peer      : 10.0.28.1
NextHop      : 10.0.28.1         Reuse time : 00h00m00s
Peer AS      : 60203             Peer Router-Id : 32.32.27.203
Local Pref   : none
Age          : 00h22m09s         Last update : 02d00h58m
FOM Present  : 738              FOM Last upd. : 2039
Number of Flaps : 2             Flags      : ud*i
Path         : 60203 65001 19855 3356 1239 22406
Applied Policy : default-damping-profile

```

 Network : 15.142.48.0/20

```

-----
Network      : 15.142.48.0/20     Peer      : 10.0.28.1
NextHop      : 10.0.28.1         Reuse time : 00h00m00s
Peer AS      : 60203             Peer Router-Id : 32.32.27.203
Local Pref   : none
Age          : 00h00m38s         Last update : 02d01h20m
FOM Present  : 2011             FOM Last upd. : 2023
Number of Flaps : 2             Flags      : ud*i
Path         : 60203 65001 19855 3356 3561 5551 1889
Applied Policy : default-damping-profile

```

 Network : 15.200.128.0/19

```

-----
Network      : 15.200.128.0/19    Peer      : 10.0.28.1
NextHop      : 10.0.28.1         Reuse time : 00h00m00s
Peer AS      : 60203             Peer Router-Id : 32.32.27.203
Local Pref   : none
Age          : 00h00m38s         Last update : 02d01h20m
FOM Present  : 2011             FOM Last upd. : 2023
Number of Flaps : 2             Flags      : ud*i
Path         : 60203 65001 19855 1299 702 1889
Applied Policy : default-damping-profile

```

 Network : 15.203.192.0/18

```

-----
Network      : 15.203.192.0/18    Peer      : 10.0.28.1

```

```

NextHop          : 10.0.28.1           Reuse time       : 00h00m00s
Peer AS          : 60203               Peer Router-Id   : 32.32.27.203
Local Pref       : none
Age              : 00h00m07s           Last update      : 02d01h20m
FOM Present      : 1018                FOM Last upd.   : 1024
Number of Flaps  : 1                   Flags            : ud*i
Path             : 60203 65001 19855 1299 702 1889
Applied Policy   : default-damping-profile

```

```
-----
*A:ALA-12#
```

```
*A:ALA-12# show router 3 bgp damping 15.203.192.0/18 detail
```

```
=====
BGP Router ID : 10.0.0.14           AS : 65206     Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best

```

```
=====
BGP Damped Routes 15.203.192.0/18
=====
```

```
Network : 15.203.192.0/18
-----
```

```

Network          : 15.203.192.0/18     Peer           : 10.0.28.1
NextHop          : 10.0.28.1           Reuse time     : 00h00m00s
Peer AS          : 60203               Peer Router-Id : 32.32.27.203
Local Pref       : none
Age              : 00h00m42s           Last update    : 02d01h20m
FOM Present      : 2003                FOM Last upd.  : 2025
Number of Flaps  : 2                   Flags          : ud*i
Path             : 60203 65001 19855 3356 702 1889
Applied Policy   : default-damping-profile

```

```
-----
Paths : 1
=====
```

```
*A: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 : 15.142.48.0/20
-----
```

```

Network          : 15.142.48.0/20     Peer           : 10.0.28.1
NextHop          : 10.0.28.1           Reuse time     : 00h29m22s
Peer AS          : 60203               Peer Router-Id : 32.32.27.203
Local Pref       : none
Age              : 00h01m28s           Last update    : 02d01h20m
FOM Present      : 2936                FOM Last upd.  : 3001
Number of Flaps  : 3                   Flags          : si
Path             : 60203 65001 19855 3356 702 1889
Applied Policy   : default-damping-profile

```

```
-----
Network : 15.200.128.0/19
-----
```

```

Network          : 15.200.128.0/19     Peer           : 10.0.28.1
NextHop          : 10.0.28.1           Reuse time     : 00h29m22s

```

VPRN Show Commands

```
Peer AS          : 60203                Peer Router-Id   : 32.32.27.203
Local Pref       : none
Age              : 00h01m28s            Last update      : 02d01h20m
FOM Present      : 2936                 FOM Last upd.   : 3001
Number of Flaps  : 3                    Flags            : si
Path             : 60203 65001 19855 3356 702 1889
Applied Policy   : default-damping-profile
```

Network : 15.203.240.0/20

```
Network          : 15.203.240.0/20      Peer             : 10.0.28.1
NextHop          : 10.0.28.1            Reuse time       : 00h29m22s
Peer AS          : 60203                Peer Router-Id   : 32.32.27.203
Local Pref       : none
Age              : 00h01m28s            Last update      : 02d01h20m
FOM Present      : 2936                 FOM Last upd.   : 3001
Number of Flaps  : 3                    Flags            : si
Path             : 60203 65001 19855 3356 702 1889
Applied Policy   : default-damping-profile
```

Network : 15.206.0.0/17

```
Network          : 15.206.0.0/17       Peer             : 10.0.28.1
NextHop          : 10.0.28.1            Reuse time       : 00h29m22s
Peer AS          : 60203                Peer Router-Id   : 32.32.27.203
Local Pref       : none
Age              : 00h01m28s            Last update      : 02d01h20m
FOM Present      : 2936                 FOM Last upd.   : 3001
Number of Flaps  : 3                    Flags            : si
Path             : 60203 65001 19855 3356 702 1889
Applied Policy   : default-damping-profile
```

*A:ALA-12#

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 **Standard and Detailed Group Output** — The following table describes the standard and detailed command output fields for a BGP group:

Sample Output

Label	Description
Group	BGP group name
Group Type	No Type — Peer type not configured. External — Peer type configured as external BGP peers. Internal — Peer type configured as internal BGP peers.
State	Disabled — The BGP peer group has been operationally disabled. Down — The BGP peer group is operationally inactive. Up — The BGP peer group is operationally active.
Peer AS	The configured or inherited peer AS for the specified peer group.
Local AS	The configured or inherited local AS for the specified peer group.
Local Address	The configured or inherited local address for originating peering for the specified peer group.
Loop Detect	The configured or inherited loop detect setting for the specified peer group.
Connect Retry	The configured or inherited connect retry timer value.
	Authentication
	None — No authentication is configured.

Label	Description (Continued)
	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 - 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.
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.

Label	Description (Continued)
Hold Time	The configured hold time setting.
Keep Alive	The configured keepalive setting.
Cluster Id	None – No cluster ID has been configured.
Client Reflect	Disabled – The BGP route reflector 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.

```
*A:ALA-12# show router 3 bgp group
=====
BGP Groups
=====
Group          : To_AS_40000
-----
Description    : Not Available
Group Type     : No Type           State          : Up
Peer AS        : 40000             Local AS       : 65206
Local Address  : n/a              Loop Detect    : Ignore
Export Policy  : direct2bgp
Hold Time      : 90                Keep Alive     : 30
Cluster Id     : None            Client Reflect : Enabled
NLRI           : Unicast          Preference     : 170

List of Peers
- 10.0.0.1     : To_Jukebox
- 10.0.0.12    : Not Available
- 10.0.0.13    : Not Available
- 10.0.0.14    : To_ALA-1
- 10.0.0.15    : To_H-215
Total Peers    : 5                Established    : 2
=====
*A:ALA-12#
```

neighbor

Syntax `neighbor [ip-address [[family family] filter1]]`
`neighbor [as-number [[family family] filter2]]`

VPRN Show Commands

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).

Note: 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 *family* — Specifies the type of routing information to be distributed by the BGP instance.

Values ipv4, vpn-ipv4

filter1 — Specifies route criteria.

Values received-routes, advertised-routes, history, suppressed, detail

filter2 — Specifies route criteria.

Values history, suppressed, detail

Output **Standard and Detailed Neighbor** — The following table describes the standard and detailed command output fields for a BGP neighbor:

Label	Description
Peer	The IP address of the configured BGP peer.
Group	The BGP peer group to which this peer is assigned.
Peer AS	The configured or inherited peer AS for the peer group.
Peer Address	The configured address for the BGP peer.
Peer Port	The TCP port number used on the far-end system.
Local AS	The configured or inherited local AS for the peer group.
Local Address	The configured or inherited local address for originating peering for the peer group.
Local Port	The TCP port number used on the local system.
Peer Type	External — Peer type configured as external BGP peers. Internal — Peer type configured as internal BGP peers.

Label	Description (Continued)
State	<p>Idle – The BGP peer is not accepting connections.</p> <p>Active – BGP is listening for and accepting TCP connections from this peer.</p> <p>Connect – BGP is attempting to establish a TCP connection from this peer.</p> <p>Open Sent – BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer.</p> <p>Open Confirm – BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION.</p> <p>Established – BGP has successfully established a peering and is exchanging routing information.</p>
Last State	<p>Idle – The BGP peer is not accepting connections.</p> <p>Active – BGP is listening for and accepting TCP connections from this peer.</p> <p>Connect – BGP is attempting to establish a TCP connection with this peer.</p> <p>Connect – BGP is attempting to establish a TCP connections from this peer.</p> <p>Open Sent – BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer.</p> <p>Open Confirm – BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION.</p> <p>Open Confirm – BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION.</p>
Last Event	<p>start – BGP has initialized the BGP neighbor.</p> <p>stop – BGP has disabled the BGP neighbor.</p> <p>open – BGP transport connection opened.</p> <p>close – BGP transport connection closed.</p> <p>openFail – BGP transport connection failed to open.</p> <p>error – BGP transport connection error.</p> <p>connectRetry – Connect retry timer expired.</p> <p>holdTime – Hold time timer expired.</p> <p>keepAlive – Keepalive timer expired.</p> <p>recvOpen – Receive an OPEN message.</p>

Label	Description (Continued)
	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.
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.

Label	Description (Continued)
	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 - 4294967295 – The maximum number of routes BGP can learn from a peer.
Hold Time	The configured hold time setting.
Keep Alive	The configured keepalive setting.
Active Hold Time	The negotiated hold time, if the BGP neighbor is in an established state.
Active Keep Alive	The negotiated keepalive time, if the BGP neighbor is in an established state.
Cluster Id	The configured route reflector cluster ID. None – No cluster ID has been configured
Client Reflect	Disabled – The BGP route reflector is configured not to reflect routes to this neighbor. Enabled – The BGP route reflector is configured to reflect routes to this neighbor.
Preference	The configured route preference value for the peer group.
Num of Flaps	The number of flaps in the neighbor connection.
Recd. Prefixes	The number of routes received from the BGP neighbor.
Active Prefixes	The number of routes received from the BGP neighbor and active in the forwarding table.
Recd. Paths	The number of unique sets of path attributes received from the BGP neighbor.
Suppressed Paths	The number of unique sets of path attributes received from the BGP neighbor and suppressed due to route damping.
Input Queue	The number of BGP messages to be processed.
Output Queue	The number of BGP messages to be transmitted.

Label	Description (Continued)
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.

Sample Output

```
*A:ALA-12# show router 3 bgp neighbor
=====
BGP Neighbor
=====
-----
Peer : 10.0.0.15          Group : To_AS_40000
-----
Peer AS      : 65205
Peer Address : 10.0.0.15      Peer Port    : 0
Local AS     : 65206
Local Address : 10.0.0.16     Local Port   : 0
Peer Type    : External
State        : Active        Last State   : Connect
Last Event   : openFail
Last Error   : Hold Timer Expire
Hold Time    : 90             Keep Alive   : 30
Active Hold Time : 0         Active Keep Alive: 0
Cluster Id   : None
Preference   : 170           Num of Flaps : 0
Recd. Prefixes : 0         Active Prefixes : 0
Recd. Paths  : 0           Suppressed Paths : 0
Input Queue  : 0           Output Queue  : 0
i/p Messages : 0           o/p Messages  : 0
i/p Octets   : 0           o/p Octets    : 0
i/p Updates  : 0           o/p Updates   : 0
Export Policy : direct2bgp
=====
*A:ALA-12#
```

```
*A:ALA-12# show router 3 bgp neighbor detail
=====
BGP Neighbor (detail)
=====
-----
Peer : 10.0.0.15          Group : To_AS_40000
-----
Peer AS      : 65205
Peer Address : 10.0.0.15      Peer Port    : 0
Local AS     : 65206
Local Address : 10.0.0.16     Local Port   : 0
```

```

Peer Type      : External
State         : Active
Last Event    : openFail
Last Error    : Hold Timer Expire
Connect Retry  : 20
Min Route Advt. : 30
Multipath     : 1
Damping       : Disabled
MED Out       : No MED Out
Next Hop Self : Disabled
Remove Private : Disabled
Prefix Limit  : No Limit
Hold Time     : 90
Active Hold Time : 0
Cluster Id    : None
Preference    : 170
Recd. Prefixes : 0
Recd. Paths   : 0
Input Queue   : 0
i/p Messages  : 0
i/p Octets    : 0
i/p Updates   : 0
Export Policy  : direct2bgp

Last State    : Connect
Local Pref.   : 100
Min AS Orig.  : 15
Multihop      : 5
Loop Detect    : Ignore
Authentication : None
AggregatorID Zero: Disabled
Passive       : Disabled

Keep Alive    : 30
Active Keep Alive: 0
Client Reflect : Enabled
Num of Flaps  : 0
Active Prefixes : 0
Suppressed Paths : 0
Output Queue  : 0
o/p Messages  : 0
o/p Octets    : 0
o/p Updates   : 0

```

```

=====
*A:ALA-12#

```

Output **Show Advertised and Received Routes Output** — The following table describes the command output fields for both the standard and detailed information for a neighbor:

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.

Label	Description (Continued)
LocalPref	BGP local preference path attribute for the route.
MED	BGP Multi-Exit Discriminator (MED) path attribute for the route.
AS Path	The BGP AS path for the route.

Sample Output

```
*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#
```

paths

Syntax paths**Context** show>router>bgp**Description** This command displays a summary of BGP path attributes.**Output** **Show Path Output** — The following table describes the command output fields for a BGP path.

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.
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.

Sample Output

```
*A:ALA-12# show router 3 bgp paths
=====
BGP Router ID : 10.0.0.14      AS : 65206   Local AS : 65206
=====
BGP Paths
```

VPRN Show Commands

```
-----  
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#
```

routes

Syntax `routes [family family] [prefix [detail | longer]]`
`routes [family family] [prefix [hunt | brief]]`
`routes [family family] [community comm-id]`
`routes [family family] [aspath-regex reg-exp1]`
`routes [family family] [ipv6-prefix[/prefix-length] [detail | longer] | [hunt [brief]]]`

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 `family family` — Specifies the type of routing information to be distributed by the BGP instance.

Values `ipv4` — Displays only those BGP peers that have the IPv4 family enable and not those capable of exchanging IP-VPN routes.
`vpn-ipv4` — Displays the BGP peers that are IP-VPN capable.
`ipv6` — Displays the BGP peers that are IPv6 capable.
`mcast-ipv4` — Displays the BGP peers that are mcast-ipv4 capable.

`prefix` — Specifies the type of routing information to display.

Values `rd[rd:]ip-address[/mask]`

<code>rd</code>	{ip-address:number1 as-number1:number2 as-number2:number3}
<code>number1</code>	1 — 65535
<code>as-number1</code>	1 — 65535
<code>number2</code>	0 — 4294967295
<code>as-number2</code>	1 — 4294967295
<code>number3</code>	0 — 65535
<code>ip-address</code>	a.b.c.d
<code>mask</code>	0 — 32

`ipv6-prefix[/prefix-length]` — Specifies the type of IPv6 routing information to display.

Values `ipv6-prefix:` `x:x:x:x:x:x:x` (eight 16-bit pieces)
`x:x:x:x:x:d.d.d`
`x:` [0 — FFFF]H
`d:` [0 — 255]D
`prefix-length` 0 — 128

`filter` — Specifies route criteria.

Values `hunt` Displays entries for the specified route in the RIB-In, RIB-Out, and RTM.
`longer` Displays the specified route and subsets of the route.
`detail` Display the longer, more detailed version of the output.

`aspath-regex "reg-exp"` — Displays all routes with an AS path matching the specified regular expression `reg-exp`.

`community comm.-id` — Displays all routes with the specified BGP community.

Values *[as-number1:comm-val1 | ext-comm | well-known-comm]*

ext-comm type: {ip-address:comm-val1 | as-number1:comm-val2 | as-number2:comm-val1 }

as-number1 0..65535

comm-val1 0..65535

type keywords: target, origin

ip-address a.b.c.d

comm-val2 0 — 4294967295

as-number2 0 — 4294967295

well-known-comm no-export, no-export-subconfed, no-advertise

Output **Show BGP Routes** — The following table describes the command output fields for BGP routes.

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured local AS setting, if not configured it is the same as the system AS.
Network	The IP prefix and mask length.
Nextthop	The BGP nextthop.
From	The advertising BGP neighbor's IP address.
Res. Nextthop	The resolved nextthop.
Local Pref.	The local preference value.
Flag	u — used s — suppressed h — history d — decayed * — valid i — igp e — egp ? — incomplete > — best
Aggregator AS	The aggregator AS value. none — No aggregator AS attributes are present.
Aggregator	The aggregator attribute value. none — no Aggregator attributes are present.
Atomic Aggr.	Atomic — The atomic aggregator flag is set.

Label	Description
	Not Atomic – The atomic aggregator flag is not set.
MED	The MED metric value. none – No MED metric is present.
Community	The BGP community attribute list.
Cluster	The route reflector cluster list.
Originator Id	The originator ID path attribute value. none – The originator ID attribute is not present.
Peer Router Id	The router ID of the advertising router.
AS-Path	The BGP AS path attribute.
VPN Imported	Displays the VPNs where a particular BGP-VPN received route has been imported and installed.

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 100.0.0.0/31 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      : 100.0.0.0/31
Nexthop      : 10.20.1.2
Route Dist.  : 10.20.1.2:1      VPN Label    : 131070
From         : 10.20.1.2
Res. Nexthop : 10.10.1.2
Local Pref.  : 100
Aggregator AS : none
Interface Name: to-sr7
Aggregator   : none
```

VPRN Show Commands

```
Atomic Aggr.   : Not Atomic           MED           : none
Community     : target:10.20.1.2:1
Cluster       : No Cluster Members
Originator Id : None                 Peer Router Id: 10.20.1.2
Flags         : Used Valid Best IGP
AS-Path       : No As-Path
VPRN Imported : 1 2 10 12
-----
RIB Out Entries
-----
Routes : 1
=====
A:SR-12#

*A:Dut-B>config>service>vprn>bgp# show router bgp routes 5.5.5.5/32 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       : 5.5.5.5/32
Nextthop     : 29.1.1.2 (VPRN 1)
Path Id      : None
From         : ::
Res. Nextthop : 29.1.1.2
Local Pref.  : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric  : None
Connector    : None
Community    : 10:2
Cluster      : No Cluster Members
Originator Id : None
Fwd Class    : 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
...
```

summary

Syntax summary [all]

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.

Output **Show BGP Summary Output** — The following table describes the command output fields for a BGP summary:

ecmp

Syntax **ecmp**

Context show>router

Description This command displays the ECMP settings for the router.

Output **Show ECMP Settings Output** — The following table describes the output fields for the router ECMP settings.

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.

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#
```

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.
- ip-int-name* — Only displays the interface information associated with the specified IP interface name.
- detail** — Displays detailed IP interface information.
- summary** — Displays summary IP interface information for the router.
- exclude-services** — Displays IP interface information, excluding IP interfaces configured for customer services. Only core network IP interfaces are displayed.

Output **Standard IP Interface Output** — The following table describes the standard output fields for an IP interface:

Label	Description
Interface-Name	The IP interface name.
Type	n/a — No IP address has been assigned to the IP interface, so the IP address type is not applicable. Pri — The IP address for the IP interface is the Primary address on the IP interface. Sec — The IP address for the IP interface is a secondary address on the IP interface.
IP-Address	The IP address and subnet mask length of the IP interface. n/a — Indicates no IP address has been assigned to the IP interface.
Adm	Down — The IP interface is administratively disabled. Up — The IP interface is administratively enabled.
Opr	Down — The IP interface is operationally disabled. Up — The IP interface is operationally enabled.
Mode	Network — The IP interface is a network/core IP interface. Service — The IP interface is a service IP interface.

Sample Output

```
*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#
```

Detailed IP Interface Output — The following table describes the detailed output fields for an IP interface.

Label	Description
If Name	The IP interface name.
Admin State	Down — The IP interface is administratively disabled. Up — The IP interface is administratively enabled.

Label	Description (Continued)
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.
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.

Label	Description (Continued)
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 (seconds)). 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.

```
*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                     Address Type   : Secondary
IGP Inhibit   : Enabled                         Broadcast Address: Host-ones
-----
Details
-----
If Index      : 2
Port Id       : 1/1/2                            If Type       : Network
Egress Filter: none                             Ingress Filter : 100
QoS Policy    : 1                               SNTP Broadcast : False
MAC Address   : 04:5d:01:01:00:02               Arp Timeout    : 14400
IP MTU        : 1500                            ICMP Mask Reply : True
Cflowd        : none

ICMP Details
Redirects     : Disabled
Unreachables  : Number - 100                     Time (seconds) - 10
TTL Expired   : Number - 100                     Time (seconds) - 10
=====
*A:ALA-12#
```

Summary IP Interface Output — The following table describes the summary output fields for the router IP interfaces.

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.

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#
```

bindings

Syntax bindings active

Context show>router>ldp

Description This command displays 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	Interface		
Group	Op	IngLbl	EgrLbl
RootAddr	EgrIf/LspId		
EgrNH			

 No Matching Entries Found
 =====

=====
 LDP In-Band-SSM IPv6 P2MP Bindings (Active)
 =====

Source

VPRN Show Commands

```

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 1.1.1.1:0)
(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/    EgrNextHop
                Peer          LspId
-----
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
1.1.1.1                               Unknw                                  --                                  131051
90.90.90.2                             1/1/6
2.2.2.2:0

104
1.1.1.1                               Unknw                                  --                                  131050
90.90.90.2                             1/1/6
2.2.2.2:0

600
1.1.1.1                               Unknw                                  --                                  131049
90.90.90.2                             1/1/6
2.2.2.2:0

700
1.1.1.1                               Unknw                                  --                                  131048
90.90.90.2                             1/1/6
2.2.2.2:0

800
1.1.1.1                               Unknw                                  --                                  131047
90.90.90.2                             1/1/6
2.2.2.2:0

900
1.1.1.1                               Unknw                                  --                                  131046
90.90.90.2                             1/1/6
2.2.2.2:0

1500
1.1.1.1                               Unknw                                  --                                  131045
90.90.90.2                             1/1/6
2.2.2.2:0

100
6.6.6.6                               Unknw                                  --                                  131044
90.90.90.2                             1/1/6
2.2.2.2:0

900
6.6.6.6                               Unknw                                  --                                  131043
90.90.90.2                             1/1/6
2.2.2.2:0

```

VPRN Show Commands

```

-----
No. of Generic IPv4 P2MP Bindings: 9
=====

LDP Generic IPv6 P2MP Bindings
=====
P2MP-Id
RootAddr                               Interface      IngLbl      EgrLbl
EgrNH                                   EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

LDP In-Band-SSM IPv4 P2MP Bindings
=====
Source
Group
RootAddr                               Interface      IngLbl      EgrLbl
EgrNH                                   EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

LDP In-Band-SSM IPv6 P2MP Bindings
=====
Source
Group
RootAddr                               Interface      IngLbl      EgrLbl
EgrNH                                   EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv4 P2MP Bindings
=====
Source
Group                                RD
RootAddr                               Interface      IngLbl      EgrLbl
EgrNH                                   EgrIf/LspId
Peer
-----
1.1.1.1
225.0.0.1                               1.1.1.1:100
3.3.3.3                                 Unknwn        --          100
60.60.60.1                              1/1/1
2.2.2.2:100

1.1.1.1
225.0.0.1                               1.1.1.1:100
3.3.3.3                                 Unknwn        --          100
60.60.60.1                              1/1/1
2.2.2.2:100

1.1.1.1
225.0.0.1                               1.1.1.1:100

```

Virtual Private Routed Network Services

```

3.3.3.3                Unknwn                --                100
60.60.60.1            1/1/1
2.2.2.2:100

```

No. of In-Band-VPN-SSM IPv4 P2MP Bindings: 3
=====

=====

```

LDP In-Band-VPN-SSM IPv6 P2MP Bindings
=====
Source
Group          RD
RootAddr      Interface      IngLbl  EgrLbl
EgrNH         EgrIf/LspId
Peer
-----
1.1.1.1
225.0.0.1     1.1.1.1:100
2000::3000    Unknwn        --      100
60.60.60.1   1/1/1
2.2.2.2:100

1.1.1.1
225.0.0.1     1.1.1.1:100
2000::3000    Unknwn        --      100
60.60.60.1   1/1/1
2.2.2.2:100

1.1.1.1
225.0.0.1     1.1.1.1:100
2000::3000    Unknwn        --      100
60.60.60.1   1/1/1
2.2.2.2:100

```

No. of In-Band-VPN-SSM IPv6 P2MP Bindings: 3
=====

=====

```

LDP Service FEC 128 Bindings
=====
Type          VCId      SDPIId  IngLbl  LMTU
Peer          SvcId     EgrLbl  RMTU
-----
?-Eth        100       R. Src  --      None
2.2.2.2:0   Ukwn     131023D 986

?-Eth        500       R. Src  --      None
2.2.2.2:0   Ukwn     131022D 1386

?-Eth        2001      R. Src  --      None
2.2.2.2:0   Ukwn     131019D 986

?-Eth        2003      R. Src  --      None
2.2.2.2:0   Ukwn     131017D 986

?-Ipipe     1800     R. Src  --      None
2.2.2.2:0   Ukwn     131014D 1486

```

No. of VC Labels: 5

VPRN Show Commands

```
=====
LDP Service FEC 129 Bindings
=====
SAII                               AGII       IngLbl     LMTU
TAII                               Type       EgrLbl     RMTU
Peer                               SvcId     SDPIId
-----
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.

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 224.1.1.1
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
=====
```

route-table

Syntax **route-table** [*ip-prefix* [*/mask*] [*longer*] | [*protocol protocol*] | [*summary*]]

Context show>router

Description This command displays the active routes in the routing table.
If no command line arguments are specified, all routes are displayed, sorted by prefix.

Parameters *ip-prefix[/mask]* — Displays routes only matching the specified *ip-prefix* and optional *mask*.

longer — Displays routes matching the *ip-prefix/mask* and routes with longer masks.

protocol *protocol* — Displays routes learned from the specified protocol.

Values *bgp, isis, local, ospf, rip, static, aggregate*

summary — Displays a route table summary information.

Output **Standard Show Route Table Output** — The following table describes the standard output fields for the route table.

Label	Description
Dest Address	The route destination address and mask.
Next Hop	The next hop IP address for the route destination.
Type	Local — The route is a local route. Remote — The route is a remote route.
Protocol	The protocol through which the route was learned.
Age	The route age in seconds for the route.
Metric	The route metric value for the route.
Pref	The route preference value for the route.
No. of Routes:	The number of routes displayed in the list.

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
```

VPRN Show Commands

```

=====
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#

```

service-prefix

Syntax **service-prefix**

Context show>router

Description This command displays service-prefix information.

Output **Show Service Prefix Output** — The following table describes the service prefix output fields.

Label	Description
IP Prefix	Displays information for the specified IP prefix.
Mask	Displays information for the specified mask length.

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#
```

static-arp

Syntax `static-arp [ip-address | ip-int-name | mac ieee-mac-addr]`

Context `show>router`

Description This command displays the router static ARP table sorted by IP address. If no options are present, all ARP entries are displayed.

Parameters *ip-address* — Only displays static ARP entries associated with the specified IP address.
ip-int-name — Only displays static ARP entries associated with the specified IP interface name.
mac ieee-mac-addr — Only displays static ARP entries associated with the specified MAC address.

Output **Static ARP Table Output** — The following table describes the output fields for the ARP table.

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.

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
12.200.1.1      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 12.200.1.1
=====
ARP Table
=====
IP Address      MAC Address      Age      Type Interface
-----
12.200.1.1      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#
```

static-route

Syntax `static-route [ip-prefix /mask] | [preference preference] | [next-hop ip-addr] [detail]`

Context show>router

Description This command displays the static entries in the routing table.

If no options are present, all static routes are displayed sorted by prefix.

Parameters *ip-prefix /mask* — Displays static routes only matching the specified *ip-prefix* and *mask*.
preference preference — Only displays static routes with the specified route preference.

Values 0 — 65535

next-hop ip-addr — Only displays static routes with the specified next hop IP address.

detail — Displays detailed information about the static route.

Output **Show Static Route Output** — The following table describes the output fields for the static route table:

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 <code>NextHop</code> for this type of route is <code>black-hole</code> . ID — The static route is an indirect route, where the <code>nextHop</code> for this type of route is the non-directly connected next hop.

Label	Description (Continued)
	NH — The route is a static route with a directly connected next hop. The <code>Nexthop</code> for this type of route is either the next hop IP address or an egress IP interface name.
Next Hop	The next hop for the static route destination.
Interface	The egress IP interface name for the static route. n/a — indicates there is no current egress interface because the static route is inactive or a 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.

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          Next Hop          Tag          Met    Pref Type Act
-----
1.2.3.4/32      10.11.25.6       0            1     5    NH   Y
ip-10.11.25.5_base_to_cpe_static
10.11.15.0/24   10.11.25.6       0            1     5    NH   Y
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          : 1.2.3.4/32
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        : enabled          State          : n/a
Target           : 10.11.18.6
Interval         : 1                Drop Count     : 3
Log              : N
CPE Host Up Time : 0d 00:00:02
CPE Echo Req Tx  : 3                CPE Echo Reply Rx : 3
CPE Up Trans     : 1                CPE Down Trans  : 0
CPE TTL          : 2
-----
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 3.3.3.3/32 detail
```

VPRN Show Commands

```
Static Route Table (Router: Base) Family: IPv4
=====
Prefix           : 3.3.3.3/32
NextHop          : n/a
Type             : Blackhole
Interface        : n/a
Prefix List      : n/a
Metric           : 1
Admin State      : Up
BFD              : disabled
CPE-check        : disabled
NextHop Type     : IP
Active           : Y
Prefix List Type : n/a
Preference       : 5
Tag              : 0
Community        : 100:33
-----
No. of Static Routes: 1
=====
```

tunnel-table

Syntax **tunnel-table** [*ip-address[/mask]*] [**protocol** *protocol* | **sdp** *sdp-id*]
tunnel-table [**summary**]

Context show>router

Description This command displays tunnel table information.

Note that auto-bind GRE tunnels are not displayed in **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 *ip-address[/mask]* — Displays the specified tunnel table's destination IP address and mask.
protocol protocol — Displays LDP protocol information.
sdp sdp-id — Displays information pertaining to the specified SDP.
summary — Displays summary tunnel table information.

Output **Show Tunnel Table Output** — The following table describes tunnel table output fields:

Label	Description
Destination	The route's destination address and mask.
Owner	Specifies the tunnel owner.
Encap	Specifies the tunnel's encapsulation type.
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.

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#

```

VPRN Clear Commands

arp

Syntax	arp
Context	clear>service>id
Description	This command clears

arp-host

Syntax	arp-host arp-host { 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] arp-host statistics [sap <i>sap-id</i> interface <i>interface-name</i>]
Context	clear>service>id
Description	This command clears ARP host data.

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	Clears DHCP statistics.

interface

Syntax	interface [<i>ip-int-name</i> <i>ip-addr</i>] [icmp]
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 *ip-int-name* | *ip-addr* — The IP interface name or IP interface address.

Default All IP interfaces.

icmp — Specifies to reset the ICMP statistics for the IP interface(s) used for ICMP rate limit.

damping

Syntax **damping** [[*ip-prefix/mask*] [**neighbor** *ip-address*]] | [**group** *name*]

Context clear>router>bgp

Description This command clears or resets the route damping information for received routes.

Parameters *ip-prefix/mask* — Clears damping information for entries that match the IP prefix and mask length.

neighbor *ip-address* — Clears damping information for entries received from the BGP neighbor.

group *name* — Clears damping information for entries received from any BGP neighbors in the peer group.

flap-statistics

Syntax **flap-statistics** [[*ip-prefix/mask*] [**neighbor** *ip-addr*]] | [**group** *group-name*] | [**regex** *reg-exp*] | [**policy** *policy-name*]

Context clear>router>bgp

Description This command clears route flap statistics.

Parameters *ip-prefix/mask* — Clears route flap statistics for entries that match the specified IP prefix and mask length.

neighbor *ip-addr* — Clears route flap statistics for entries received from the specified BGP neighbor.

group *group-name* — Clears route flap statistics for entries received from any BGP neighbors in the specified peer group.

regex *reg-exp* — Clears route flap statistics for all entries which have the regular expression and the AS path that matches the regular expression.

policy *policy-name* — Clears route flap statistics for entries that match the specified route policy.

neighbor

Syntax **neighbor** {*ip-addr* | **as** *as-number* | **external** | **all**} [**soft** | **soft-inbound** | **statistics**]

Context clear>router>bgp

Description This command resets the specified BGP peer or peers. This can cause existing BGP connections to be shutdown and restarted.

Parameters *ip-addr* — Resets the BGP neighbor with the specified IP address.

VPRN Show Commands

as *as-number* — Resets all BGP neighbors with the specified peer AS.

external — Resets all EBGp neighbors.

all — Resets all BGP neighbors.

soft — The specified BGP 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.

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.

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 — 2147483648

sap

Syntax **sap** *sap-id* {**all** | **counters** | **stp**}

Context clear>service>statistics

Description Clears SAP statistics for a SAP.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See [Common CLI Command Descriptions on page 849](#) for command syntax.

dhcp

Syntax **dhcp**

Context clear>router>dhcp

Description This command enables the context to clear DHCP parameters.

lease-state

Syntax **lease-state**
lease-state ip-address *ip-address*
lease-state mac *ieee-address*
lease-state sap *sap-id*
lease-state sdp *sdp-id:vc-id*

Context clear>service>id>dhcp

Description Clears DHCP lease state information for this service.

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).

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.

sap-id — Specifies the physical port identifier portion of the SAP definition. See [Common CLI Command Descriptions on page 849](#) for command syntax.

sdp-id — The SDP ID to be cleared.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID to be cleared.

Values 1 — 4294967295

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 in length.

Values 1 — 2147483648

spoke-sdp

Syntax **spoke-sdp** *sdp-id:vc-id ingress-vc-label*

Context clear>service>id

Description This command clears and resets the spoke SDP bindings for the service.

VPRN Show Commands

Parameters *sdp-id* — The spoke SDP ID to be reset.
Values 1 — 17407
vc-id — The virtual circuit ID on the SDP ID to be reset.
Values 1 — 4294967295

sdp

Syntax **sdp** *sdp-id* **keep-alive**
Context clear>service>statistics
Description This command clears keepalive statistics associated with the SDP ID.
Parameters *sdp-id* — The SDP ID for which to clear keepalive statistics.
Values 1 — 17407

counters

Syntax **counters**
Context clear>service>statistics>id
Description Clears all traffic queue counters associated with the service ID.

spoke-sdp

Syntax **spoke-sdp** *sdp-id[:vc-id]* {**all** | **counters** | **stp**}
Context clear>service>statistics>id
Description This command clears statistics for the spoke SDP bound to the service.
Parameters *sdp-id* — The spoke SDP ID for which to clear statistics.
Values 1 — 17407
vc-id — The virtual circuit ID on the SDP ID to be reset.
Values 1 — 4294967295
all — Clears all queue statistics and STP statistics associated with the SDP.
counters — Clears all queue statistics associated with the SDP.
stp — Clears all STP statistics associated with the SDP.

stp

Syntax **stp**

Context clear>service>statistics>id

Description Clears all spanning tree statistics for the service ID.

VPRN Debug Commands

id

Syntax [no] id *service-id*

Context debug>service

Description This command debugs commands for a specific service.
The **no** form of the command disables debugging.

Parameters *service-id* — The ID that uniquely identifies 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.

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 *ip-address*

Context debug>service>id>host-connectivity-verify

Description This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular IP address.

Parameters *ip-address* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

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.

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. See Common CLI Command Descriptions on page 849 for command syntax.

sap

Syntax	[no] sap <i>sap-id</i>
Context	debug>service>id debug>service>stp
Description	This command enables STP debugging for a specific SAP. The no form of the command disables debugging.
Parameters	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 849 for command syntax.

sdp

Syntax	[no] sdp <i>sdp-id:vc-id</i>
Context	debug>service>id debug>service>stp
Description	This command enables STP debugging for a specific SDP. The no form of the command disables debugging.

event-type

Syntax [no] event-type {config-change | svc-oper-status-change | sap-oper-status-change | sdpbind-oper-status-change}

Context debug>service>id

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}

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.

igmp

Syntax [no] igmp

Context debug>router

Description This command enables debugging for IGMP.
The **no** form of the command disables debugging.

interface

Syntax [no] interface [*ip-int-name* | *ip-address*]

Context debug>router>igmp

Description This command enables debugging on the IGMP interface.
The **no** form of the command disables debugging.

Parameters *ip-int-name* — Only displays the information associated with the specified IP interface name.
ip-address — Only displays the information associated with the specified IP address.

Sample Output

```
A:FA# debug router 100 igmp interface
A:FA#
A:FA# show debug
debug
    router "100"
        igmp
            interface
        exit
    exit
exit
*A:FA#
38397 2007/02/01 11:46:40.94 UTC MINOR: DEBUG #2001 vprn100 IGMP[85]
"IGMP[85]: INTF
Querier Timer expired on i/f 2"

38398 2007/02/01 11:46:40.94 UTC MINOR: DEBUG #2001 vprn100 IGMP[85]
"IGMP[85]: INTF
Sending query on i/f 2 to 0.0.0.0"

38399 2007/02/01 11:46:40.94 UTC MINOR: DEBUG #2001 vprn100 IGMP[85]
```

```

"IGMP[85]: INTF
Interface 2 already UP, ignoring event"

38400 2007/02/01 11:46:41.64 UTC MINOR: DEBUG #2001 vprn100 IGMP[85]
"IGMP[85]: INTF
Process received group rec MODE_IS_EXCL for i/f 2 group 225.1.1.1 in mode EXCLUD
E. Num srcls 0"

38401 2007/02/01 11:46:41.64 UTC MINOR: DEBUG #2001 vprn100 IGMP[85]
"IGMP[85]: INTF
Process received group rec MODE_IS_EXCL for i/f 2 group 225.1.1.2 in mode EXCLUD
E. Num srcls 0"

38402 2007/02/01 11:46:41.64 UTC MINOR: DEBUG #2001 vprn100 IGMP[85]
"IGMP[85]: INTF
Process received group rec MODE_IS_EXCL for i/f 2 group 225.1.1.3 in mode EXCLUD
E. Num srcls 0"

38403 2007/02/01 11:46:41.64 UTC MINOR: DEBUG #2001 vprn100 IGMP[85]
"IGMP[85]: INTF
Process received group rec MODE_IS_EXCL for i/f 2 group 225.1.1.4 in mode EXCLUD
E. Num srcls 0"

38404 2007/02/01 11:46:41.64 UTC MINOR: DEBUG #2001 vprn100 IGMP[85]
"IGMP[85]: INTF
Process received group rec MODE_IS_EXCL for i/f 2 group 225.1.1.5 in mode EXCLUD
E. Num srcls 0"

38405 2007/02/01 11:46:48.93 UTC MINOR: DEBUG #2001 vprn100 IGMP[85]
"IGMP[85]: INTF
Process received group rec MODE_IS_EXCL for i/f 2 group 225.1.1.1 in mode EXCLUD
E. Num srcls 0"
38408 2007/02/01 11:46:48.93 UTC MINOR: DEBUG #2001 vprn100 IGMP[85]
"IGMP[85]: INTF
Process received group rec MODE_IS_EXCL for i/f 2 group 225.1.1.4 in mode EXCLUD
E. Num srcls 0"

38409 2007/02/01 11:46:48.93 UTC MINOR: DEBUG #2001 vprn100 IGMP[85]
"IGMP[85]: INTF
Process received group rec MODE_IS_EXCL for i/f 2 group 225.1.1.5 in mode EXCLUD
E. Num srcls 0"

38410 2007/02/01 11:46:48.93 UTC MINOR: DEBUG #2001 vprn100 IGMP[85]
"IGMP[85]: INTF
Interface 2 already UP, ignoring event"
A:FA#

```

mcs

Syntax `[no] mcs [ip-int-name]`

Context `debug>router>igmp`

Description This command enables debugging for IGMP MCS.
The **no** form of the command disables debugging.

Parameters *ip-int-name* — Only displays the information associated with the specified IP interface name.

Sample Output

VPRN Debug Commands

```
A:BA# debug router 100 igmp mcs
A:BA# show debug
debug
  router "100"
  igmp
  mcs
  exit
  exit
exit
A:BA#
```

misc

Syntax [no] misc

Context debug>router>igmp

Description This command enables debugging for IGMP miscellaneous. The **no** form of the command disables debugging.

Sample Output

```
A:BA# debug router 100 igmp misc
A:BA# show debug
debug
  router "100"
  igmp
  misc
  exit
  exit
exit
A:BA#
```

packet

Syntax [no] packet [*query*|*v1-report*|*v2-report*|*v3-report*|*v2-leave*] [*ip-int-name* | *ip-address*]

Context debug>router>igmp

Description This command enables debugging for IGMP packets. The **no** form of the command disables debugging.

Parameters *query v1/v2/v3-report, v2-leave* — Select the type of packet to debug.
ip-int-name — Only displays the information associated with the specified IP interface name.
ip-address — Only displays the information associated with the specified IP address.

Sample Output

```
A:BA# debug router 100 igmp packet
A:BA#
A:BA# show debug
debug
```

```

router "100"
  igmp
    packet
  exit
exit
exit
5 2006/09/03 22:20:05.73 UTC MINOR: DEBUG #2001 vprn100 IGMP[2]
"IGMP[2]: TX-PKT
[000 18:25:24.480] ifId:2 ifName:IGMP_to_CE IGMP V3 PDU: 11.1.1.1 -> 224.0.0.1 p
duLen 12
  Type: QUERY maxrespCode 0xa checksum 0xec78
  GroupAddr: 0.0.0.0
  S bit 0, QRV 2, QQIC 125, NumSources 0
  Source Address List:
"
6 2006/09/03 22:20:05.96 UTC MINOR: DEBUG #2001 vprn100 IGMP[2]
"IGMP[2]: RX-PKT
[000 18:25:24.710] ifId:2 ifName:IGMP_to_CE IGMP V3 PDU: 11.1.1.20 -> 224.0.0.22
pduLen 48
  Type: V3 REPORT maxrespCode 0x0 checksum 0x5fe2
  Num Group Records: 4
  Group Record 0
  Type: CHG_TO_EXCL, AuxDataLen 0, Num Sources 0
  Mcast Addr: 225.1.1.1
  Source Address List
  Group Record 1
  Type: CHG_TO_EXCL, AuxDataLen 0, Num Sources 0
  Mcast Addr: 225.1.1.2
  Source Address List
  Group Record 2
  Type: CHG_TO_EXCL, AuxDataLen 0, Num Sources 0
  Mcast Addr: 225.1.1.3
  Source Address List
  Group Record 3
  Type: CHG_TO_EXCL, AuxDataLen 0, Num Sources 0
  Mcast Addr: 225.1.1.4
  Source Address List
A:BA#
*A:BA# no debug
Trace disabled for all existing and future clients
*A:BA# show debug
debug
exit

```


Common CLI Command Descriptions

In This Chapter

This section provides information about common Command Line Interface (CLI) syntax and command usage.

Topics in this chapter include:

- [SAP syntax on page 850](#)

Common Service Commands

sap

Syntax [no] sap *sap-id*

Description This command specifies the physical port identifier portion of the SAP definition.

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

The *sap-id* can be configured in one of the following formats:

Type	Syntax	Example
port-id	<i>slot/mda/port[.channel]</i>	1/1/5
null	[<i>port-id</i> <i>lag-id</i>]	<i>port-id</i> : 1/1/3 <i>lag-id</i> : lag-3
dot1q	[<i>port-id</i> <i>lag-id</i>]:qtag1	<i>port-id</i> :qtag1: 1/1/3:100 <i>lag-id</i> :qtag1:lag-3:102
qinq	[<i>port-id</i> <i>lag-id</i>]:qtag1.qtag2	<i>port-id</i> :qtag1.qtag2: 1/1/3:100.10 <i>lag-id</i> :qtag1.qtag2: lag-10:

The values depends on the encapsulation type configured for the interface. The following table describes the allowed values for the port and encapsulation types.

Port Type	Encap-Type	Allowed Values	Comments
Ethernet	Null	0	The SAP is identified by the port.
Ethernet	Dot1q	0 — 4094	The SAP is identified by the 802.1Q tag on the port. Note that a 0 qtag1 value also accepts untagged packets on the dot1q port.
Ethernet	QinQ	qtag1: 0 — 4094 qtag2: 0 — 4094	The SAP is identified by two 802.1Q tags on the port. Note that a 0 qtag1 value also accepts untagged packets on the Dot1q port.
SONET/SDH	IPCP	-	The SAP is identified by the channel. No BCP is deployed and all traffic is IP.
SONET/SDH	BCP-Null	0	The SAP is identified with a single service on the channel. Tags are assumed to be part of the customer packet and not a service delimiter.

SONET/SDH BCP-Dot1q 0 — 4094

The SAP is identified by the 802.1Q tag on the channel.

Standards and Protocol Support



Note: The information presented is subject to change without notice.

Alcatel-Lucent assumes no responsibility for inaccuracies contained herein.

ANCP/L2CP

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

draft-ietf-ancp-protocol-02, *Protocol for Access Node Control Mechanism in Broadband Networks*

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*

BGP

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 5575, *Dissemination of Flow Specification Rules*
- RFC 5668, *4-Octet AS Specific BGP Extended Community*
- draft-ietf-idr-add-paths-04, *Advertisement of Multiple Paths in BGP*
- draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*

Circuit Emulation

- MEF-8, *Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks, October 2004*
- 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, *Ethernet Linear Protection Switching*
- ITU-T G.8032, *Ethernet Ring Protection Switching*
- ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

EVPN

- RFC7432, *BGP MPLS-Based Ethernet VPN*
- draft-ietf-bess-evpn-overlay-01, *A Network Virtualization Overlay Solution using EVPN*
- draft-ietf-bess-evpn-prefix-advertisement-01, *IP Prefix Advertisement in EVPN*

draft-ietf-bess-evpn-vpls-seamless-integ-00, *(PBB-)EVPN Seamless Integration with (PBB-)VPLS*
 draft-ietf-l2vpn-pbb-evpn-10, *Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)*
 draft-snr-bess-evpn-proxy-arp-nd-00, *Proxy-ARP/ND function in EVPN networks*

Fast Reroute

RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*
 RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*
 draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*
 draft-kastran-moftr-02, *Multicast only Fast Re-Route*

Frame Relay

ANSI T1.617 Annex D, *DSSI - 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*

IP — General

RFC 768, *User Datagram Protocol*
 RFC 793, *Transmission Control Protocol*
 RFC 854, *TELNET Protocol Specifications*

RFC 951, *Bootstrap Protocol (BOOTP)*
 RFC 1034, *Domain Names - Concepts and Facilities*
 RFC 1035, *Domain Names - Implementation and Specification*
 RFC 1350, *The TFTP Protocol (revision 2)*
 RFC 1534, *Interoperation between DHCP and BOOTP*
 RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*
 RFC 2131, *Dynamic Host Configuration Protocol*
 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 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 3046, *DHCP Relay Agent Information Option (Option 82)*
 RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*
 RFC 3596, *DNS Extensions to Support IP version 6*
 RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*
 RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*
 RFC 4251, *The Secure Shell (SSH) Protocol Architecture*
 RFC 4254, *The Secure Shell (SSH) Connection Protocol*
 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*
draft-grant-tacacs-02, *The TACACS+ Protocol*
draft-ietf-vrrp-unified-spec-02, *Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6*

IP — Multicast

RFC 1112, *Host Extensions for IP Multicasting*
RFC 2236, *Internet Group Management Protocol, Version 2*
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 4601, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*
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 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*
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 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*
draft-dolganow-l3vpn-mvpn-expl-track-00, *Explicit tracking in MPLS/BGP IP VPNs*

IP — Version 4

RFC 791, *Internet Protocol*
RFC 792, *Internet Control Message Protocol*
RFC 826, *An Ethernet Address Resolution Protocol*
RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*
RFC 1812, *Requirements for IPv4 Routers*
RFC 2401, *Security Architecture for Internet Protocol*

RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*

IP — Version 6

RFC 1981, *Path MTU Discovery for IP version 6*

RFC 2460, *Internet Protocol, Version 6 (IPv6) Specification*

RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*

RFC 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*

RFC 3587, *IPv6 Global Unicast Address Format*

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 Only)*

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 5952, *A Recommendation for IPv6 Address Text Representation*

RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*

RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*

IPsec

RFC 2401, *Security Architecture for the Internet Protocol*

RFC 2406, *IP Encapsulating Security Payload (ESP)*

RFC 2409, *The Internet Key Exchange (IKE)*

RFC 2560, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*

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 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 5996, *Internet Key Exchange Protocol Version 2 (IKEv2)*

RFC 5998, *An Extension for EAP-Only Authentication in IKEv2*

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*

draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*

IS-IS

ISO/IEC 10589:2002, Second Edition, Nov. 2002, *Intermediate system to Intermediate system intra-domain routeing 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*

Standards and Protocol Support

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*
draft-ietf-isis-mi-02, *IS-IS Multi-Instance*

draft-ietf-isis-segment-routing-extensions-04, *IS-IS Extensions for Segment Routing*
draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*

Management

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*
SFLOW-MIB, *sFlow MIB Version 1.3 (Draft 5)*
RFC 1157, *A Simple Network Management Protocol (SNMP)*
RFC 1215, *A Convention for Defining Traps for use with the SNMP*
RFC 1724, *RIP Version 2 MIB Extension*
RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIPv2*
RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIPv2*
RFC 2138, *Remote Authentication Dial In User Service (RADIUS)*
RFC 2206, *RSVP Management Information Base using SMIPv2*
RFC 2213, *Integrated Services Management Information Base using SMIPv2*
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 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 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 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 4292, *IP Forwarding Table MIB*
- RFC 4293, *Management Information Base for the Internet Protocol (IP)*
- 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 6241, *Network Configuration Protocol (NETCONF)*

- RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*
- draft-ietf-snmv3-update-mib-05, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*
- draft-ietf-idr-bgp4-mib-05, *Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4)*
- 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 SMIPv2*
- 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*

MPLS — General

- RFC 3031, *Multiprotocol Label Switching Architecture*
- RFC 3032, *MPLS Label Stack Encoding*
- RFC 3443, *Time To Live (TTL) Processing in Multiprotocol 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*

MPLS — GMPLS

- 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*
- draft-ietf-ccamp-rsvp-te-srlg-collect-04, *RSVP-TE Extensions for Collecting SRLG Information*

MPLS — LDP

- 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 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*
- RFC 6826, *Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*
- draft-ietf-mpls-ldp-ip-pw-capability-09, *Controlling State Advertisements Of Non-negotiated LDP Applications*
- draft-ietf-mpls-ldp-ipv6-15, *Updates to LDP for IPv6*

draft-pdutta-mpls-ldp-adj-capability-00, *LDP Adjacency Capabilities*
 draft-pdutta-mpls-ldp-v2-00, *LDP Version 2*
 draft-pdutta-mpls-multi-ldp-instance-00, *Multiple LDP Instances*
 draft-pdutta-mpls-tldp-hello-reduce-04, *Targeted LDP Hello Reduction*

MPLS — 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 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*

MPLS — OAM

RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*
 RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*
 RFC 6425, *Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping*

MPLS — RSVP-TE

RFC 2702, *Requirements for Traffic Engineering over MPLS*

RFC 2747, *RSVP Cryptographic Authentication*
 RFC 2961, *RSVP Refresh Overhead Reduction Extensions*
 RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value*
 RFC 3209, *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*

Standards and Protocol Support

RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*

draft-newton-mpls-te-dynamic-overbooking-00, *A Diffserv-TE Implementation Model to dynamically change booking factors during failure events*

NAT

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 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*

OpenFlow

ONF *OpenFlow Switch Specification Version 1.3.1* (OpenFlow-hybrid switches)

OSPF

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 4970, *Extensions to OSPF for Advertising Optional Router Capabilities*

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*

draft-ietf-ospf-prefix-link-attr-06, *OSPFv2 Prefix/Link Attribute Advertisement*

draft-ietf-ospf-segment-routing-extensions-04, *OSPF Extensions for Segment Routing*

Policy Management and Credit Control

3GPP TS 29.212, *Policy and Charging Control (PCC) over Gx/Sd Reference Point (Release 11 and Release 12) Gx support as it applies to wireline environment (BNG)*

RFC 3588, *Diameter Base Protocol*

RFC 4006, *Diameter Credit-Control Application*

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 2661, *Layer Two Tunneling Protocol "L2TP"*
- RFC 2686, *The Multi-Class Extension to Multi-Link PPP*
- RFC 2878, *PPP Bridging Control Protocol (BCP)*
- RFC 4951, *Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"*
- RFC 5072, *IP Version 6 over PPP*

Pseudowire

- MFA Forum 9.0.0, *The Use of Virtual trunks for ATM/MPLS Control Plane 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*
- 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*
- draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS Interworking*

Quality of Service

- 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 3260, *New Terminology and Clarifications for Diffserv*
- RFC 2598, *An Expedited Forwarding PHB*
- RFC 3140, *Per Hop Behavior Identification Codes*

RIP

- RFC 1058, *Routing Information Protocol*
- RFC 2080, *RIPng for IPv6*
- RFC 2082, *RIP-2 MD5 Authentication*
- RFC 2453, *RIP Version 2*

SONET/SDH

- ITU-G.841, *Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1, issued in July 2002*

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*
- 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*

Voice and Video Performance

- 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)*

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)*

Customer Documentation and Product Support



Customer Documentation

<http://documentation.alcatel-lucent.com>



Technical Support

<http://support.alcatel-lucent.com>



Documentation Feedback

documentation.feedback@alcatel-lucent.com

