7750 SR OS Services Guide

Software Version: 7750 SR OS 5.0
February 2007
Document Part Number: 93-0076-03-01

Table of Contents

Virtual Private LAN Service

List of Figures

Services Overview

VLL Services

Virtual Private LAN Service

Internet Enhanced Service

Virtual Private Routed Network Service

Versatile Service Module

Mirror Services

OAM and SAA

List of Tables

Versatile Service Module

Mirror Services

OAM and SAA

Getting Started

In This Chapter

This book provides process flow information to configure provision services.

Alcatel-Lucent 7750 SR-Series Services Configuration Process

Table 1 lists the tasks necessary to configure subscriber services and configure mirroring. This guide is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

Table 1: 7750 SR Configuration Process

Area	Task	Chapter
Subscriber services	Configure subscriber services	
	Global entities	Configuring Global Service Entities with CLI on page 55
	VLL services:	
	Epipe service	Ethernet Pipe (Epipe) Services on page 110
	Apipe service	ATM VLL (Apipe) Services on page 113
	Fpipe service	Frame Relay VLL (Fpipe) Services on page 118
	Ipipe	IP Interworking VLL (Ipipe) Services on page 122
	VPLS service	Virtual Private LAN Service on page 323
	IES service	Internet Enhanced Service on page 661

Table 1: 7750 SR Configuration Process (Continued)

Area	Task	Chapter
	VPRN service	
Service cross connection	Cross connect aggregation groups	
Diagnostics/Service verification	Mirroring	Mirror Services on page 1295
	OAM	OAM and SAA on page 1367
Reference	List of IEEE, IETF, and other proprietary entities.	Standards and Protocol Support on page 1471

Preface

About This Guide

This guide describes subscriber services, and mirroring support provided by the 7750 SR OS and presents examples to configure and implement various protocols and services.

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

Audience

This manual is intended for network administrators who are responsible for configuring the 7750 SR routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Protocols, standards, and services described in this manual include the following:

- CLI concepts
- Subscriber services
- Service mirroring
- Operation, Administration and Maintenance (OAM) operations

List of Technical Publications

The 7750 SR documentation set is composed of the following books:

- 7750 SR OS Basic System Configuration Guide

  This guide describes basic system configurations and operations.

- 7750 SR OS System Management Guide

  This guide describes system security and access configurations as well as event logging and accounting logs.

- 7750 SR OS Interface Configuration Guide

  This guide describes card, Media Dependent Adapter (MDA), and port provisioning.

- 7750 SR OS Router Configuration Guide

  This guide describes logical IP routing interfaces and associated attributes such as an IP address, port, link aggregation group (LAG) as well as IP and MAC-based filtering, VRRP, and Cflowd.

- 7750 SR OS Routing Protocols Guide

  This guide provides an overview of routing concepts and provides configuration examples for RIP, OSPF, IS-IS, Multicast, BGP, and route policies.

- 7750 SR OS MPLS Guide

  This guide describes how to configure Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP).

- 7750 SR OS Services Guide

  This guide describes how to configure service parameters such as service distribution points (SDPs), customer information, user services, service mirroring and Operations, Administration and Management (OAM) tools.

- 7750 SR OS Triple Play Guide

  This guide describes Triple Play services and support provided by the 7750 SR and presents examples to configure and implement various protocols and services.

- 7750 SR Quality of Service Guide

  This guide describes how to configure Quality of Service (QoS) policy management.

Technical Support

If you purchased a service agreement for your 7750 SR 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, contact your welcome center.

Web: http://www1.alcatel-lucent.com/comps/pages/carrier_support.jhtml

Services Command Reference

In This Chapter

This chapter provides the command reference trees for the 7750 SR services.

Topics include:

- Global Services Commands
  - → Customer Commands on page 75
  - → SDP Commands on page 76
- Service Configuration Commands
  - → Apipe Service Configuration Commands on page 221
  - → Epipe Service Configuration Commands on page 224
  - → Fpipe Service Configuration Commands on page 227
  - → Ipipe Service Configuration Commands on page 229
  - → VPLS Service Configuration Commands on page 441
  - → IES Service Configuration Commands on page 705
  - → VPRN Service Configuration Commands on page 925

Services Overview

In This Section

This section provides an overview of the SR-Series subscriber services, service model and service entities. Additional details on the individual subscriber services can be found in subsequent chapters.

Topics in this section include:

Introduction

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

Services can provide Layer 2/bridged service or Layer3/IP routed connectivity between a service access point (SAP) on one SR-Series router and another service access point (a SAP is where traffic enters and exits the service) on the same (local) or another SR-Series router (distributed). A distributed service spans more than one router.

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

Service Types

The SR-Series offers the following types of subscriber services which are described in more detail in the referenced chapters:

- Virtual Leased Line (VLL) services:
  - → Ethernet pipe (Epipe) — A Layer 2 point-to-point VLL service for Ethernet frames. See Ethernet Pipe (Epipe) Services on page 110.
  - → ATM VLL (Apipe) — A point-to-point ATM service between users connected to 7750 nodes on an IP/MPLS network. See ATM VLL (Apipe) Service Overview on page 170.
  - → Frame-Relay (Fpipe) — A point-to-point Frame Relay service between users connected to 7750 nodes on the IP/MPLS network. See Frame Relay VLL (Fpipe) Service Overview on page 172.
  - → IP Pipe (Ipipe) — Provides IP connectivity between a host attached to a point-to-point access circuit (FR, ATM, PPP) with routed IPv4 encapsulation and a host attached to an Ethernet interface. See IP Interworking VLL (Ipipe) Overview on page 178.
- Virtual Private LAN Service (VPLS) — A Layer 2 multipoint-to-multipoint VPN. See Virtual Private LAN Service on page 323. VPLS includes Hierarchical VPLS (H-VPLS) which is an enhancement of VPLS which extends Martini-style signaled or static virtual circuit labeling outside the fully meshed VPLS core.

- Internet Enhanced Service (IES) — A direct Internet access service where the customer is assigned an IP interface for Internet connectivity. See Internet Enhanced Service on page 661.
- Virtual Private Routed Network (VPRN) — A layer 3 IP multipoint-to-multipoint VPN service as defined in RFC 2547bis. See Virtual Private Routed Network Service on page 855.

Service Policies

Common to all SR-Series connectivity services are policies that are assigned to the service. Policies are defined at a global level and then applied to a service on the router. Policies are used to define SR-Series service enhancements. The types of policies that are common to all SR-Series connectivity services are:

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

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

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

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

- Scheduler policies define the hierarchy and operating parameters for virtual schedulers. Schedulers are divided into groups based on the tier each scheduler is created under. A tier is used to give structure to the schedulers within a policy and define rules for parent scheduler associations.

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

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

Multipoint Shared Queuing

Multipoint shared queuing is supported to minimize the number of multipoint queues created for ingress VPLS, IES or VPRN SAPs or ingress subscriber SLA profiles. Normally, ingress multipoint packets are handled by multipoint queues created for each SAP or subscriber SLA profile instance. In some instances, the number of SAPs or SLA profile instances are sufficient for the in use multipoint queues to represent many thousands of queues on an ingress forwarding plane. If multipoint shared queuing is enabled for the SAPs or SLA profile instances on the forwarding plane, the multipoint queues are not created. Instead, the ingress multipoint packets are handled by the unicast queue mapped to the forwarding class of the multipoint packet.

Functionally, multipoint shared queuing is a superset of shared queuing. With shared queuing on a SAP or SLA profile instance, only unicast packets are processed twice, once for the initial service level queuing and a second time for switch fabric destination queuing. Shared queuing does not affect multipoint packet handling. Multipoint packet handling in normal (service queuing) is the same as shared queuing. When multipoint shared queuing is enabled, shared queuing for unicast packets is automatically enabled.

Ingress Queuing Modes of Operation

Three modes of ingress SAP queuing are supported for multipoint services (IES, VPLS and VPRN); service, shared, and multipoint shared. The same ingress queuing options are available for IES and VPLS subscriber SLA profile instance queuing.

Ingress Service Queuing

Normal or service queuing is the default mode of operation for SAP ingress queuing. Service queuing preserves ingress forwarding bandwidth by allowing a service queue defined in an ingress SAP QoS policy to be represented by a group of hardware queues. A hardware queue is created for each switch fabric destination to which the logical service queue must forward packets. For a VPLS SAP with two ingress unicast service queues, two hardware queues are used for each destination forwarding engine the VPLS SAP is forwarding to. If three switch fabric destinations are involved, six queues are allocated (2 unicast service queues multiplied by 3 destination forwarding complexes equals six hardware queues). Figure 1 demonstrates unicast hardware queue expansion. Service multipoint queues in the ingress SAP QoS policy are not expanded to multiple hardware queues, each service multipoint queue defined on the SAP equates to a single hardware queue to the switch fabric.

When multiple hardware queues represent a single logical service queue, the system automatically monitors the offered load and forwarding rate of each hardware queue. Based on the monitored state of each hardware queue, the system imposes an individual CIR and PIR rate for each queue that provides an overall aggregate CIR and PIR reflective of what is provisioned on the service queue.

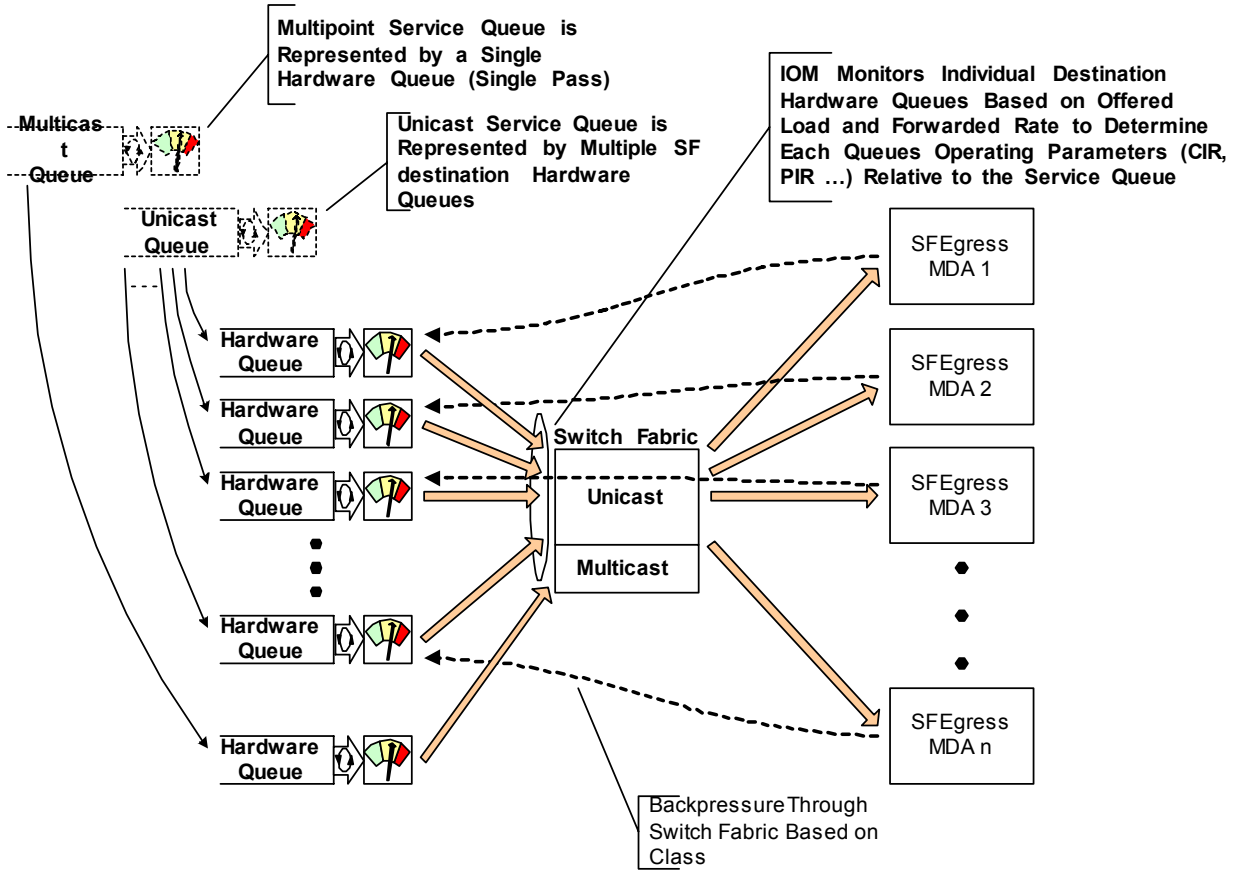

Figure 1: Unicast Service Queue Mapping to Multiple Destination Based Hardware Queues

Ingress Shared Queuing

To avoid the hardware queue expansion issues associated with normal service based queuing, the system allows an ingress logical service queue to map to a single hardware queue when shared queuing is enabled. Shared queuing uses two passes through the ingress forwarding plane to separate ingress per service queuing from the destination switch fabric queuing. In the case of shared queuing, ingress unicast service queues are created one-for-one relative to hardware queues. Each hardware queue representing a service queue is mapped to a special destination in the traffic manager that 'forwards' the packet back to the ingress forwarding plane allowing a second pass through the traffic manager. In the second pass, the packet is placed into a 'shared' queue for the destination forwarding plane. The shared queues are used by all services configured for shared queuing.

When the first SAP or SLA profile instance is configured for shared queuing on an ingress forwarding plane, the system allocates eight hardware queues per available destination forwarding plane, one queue per forwarding class. (Twenty four hardware queues are also allocated for

multipoint shared traffic, but that is discussed in the following section.) The shared queue parameters that define the relative operation of the forwarding class queues are derived from the Shared Queue policy defined in the QoS CLI node. Figure 2 demonstrates shared unicast queuing. SAP or SLA profile instance multipoint queuing is not affected by enabling shared queuing. Multipoint queues are still created as defined in the ingress SAP QoS policy and ingress multipoint packets only traverse the ingress forwarding plane a single time.

Enabling shared queuing may affect ingress performance due to double packet processing through the service and shared queues.

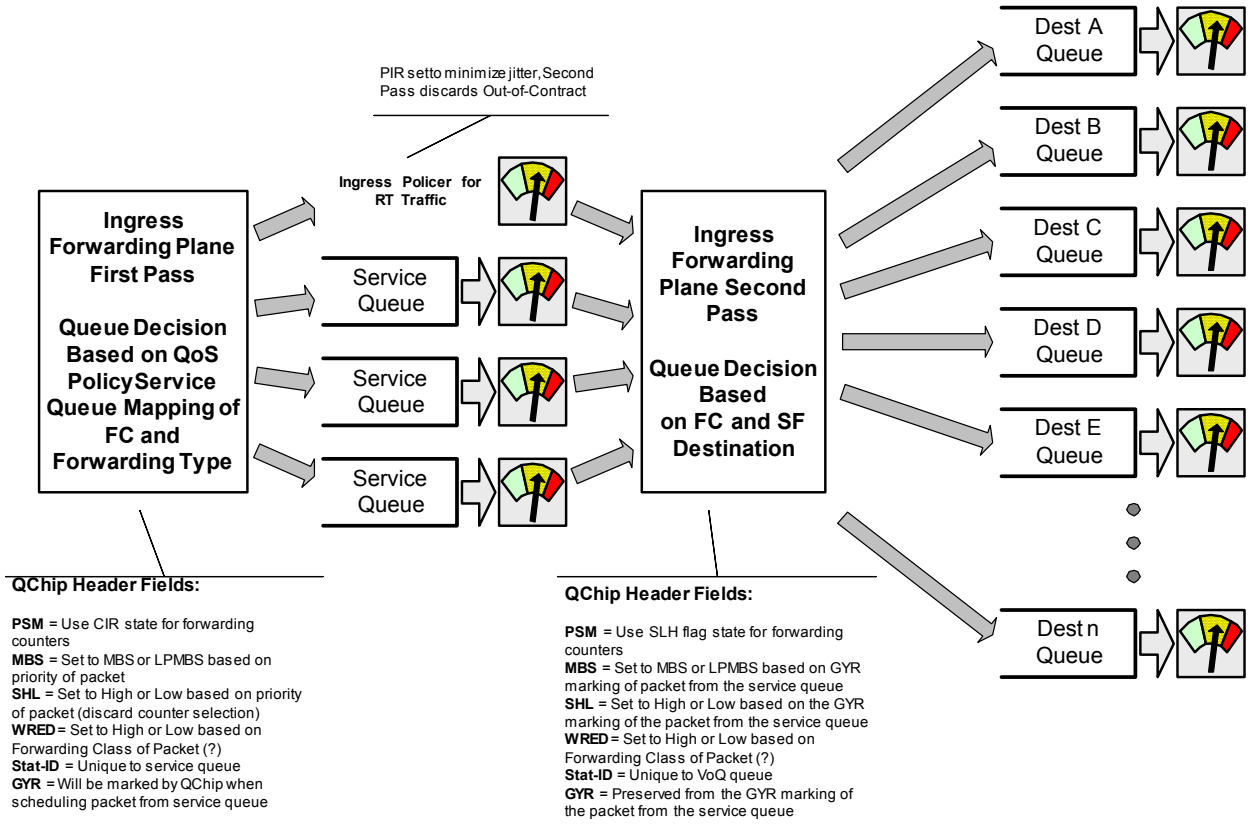

Figure 2: Unicast Service Queuing With Shared Queuing Enabled

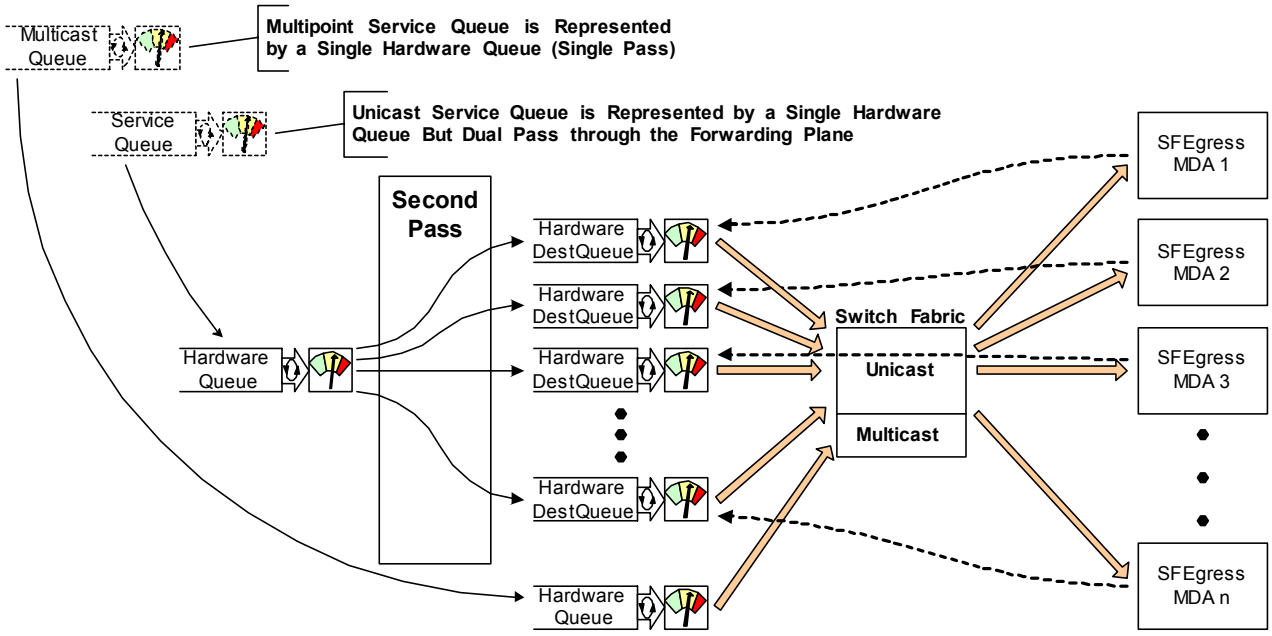

Figure 3: Multipoint Queue Behavior with Shared Queuing Enabled

Ingress Multipoint Shared Queuing

Ingress multipoint shared queuing is a variation to the unicast shared queuing defined in Ingress Shared Queuing on page 30. Ingress unicast service queues are mapped one-for-one with hardware queues and unicast packets traverse the ingress forwarding plane twice. In addition to the above, the multipoint queues defined in the ingress SAP QoS policy are not created. Instead, multipoint packets (broadcast, multicast and unknown unicast destined) are treated to the same dual pass ingress forwarding plane processing as unicast packets. In the first pass, the forwarding plane uses the unicast queue mappings for each forwarding plane. The second pass uses the multipoint shared queues to forward the packet to the switch fabric for special replication to all egress forwarding planes that need to process the packet.

The benefit of defining multipoint shared queuing is the savings of the multipoint queues per service. By using the unicast queues in the first pass and then the aggregate shared queues in the second pass, per service multipoint queues are not required. The predominate scenario where multipoint shared queuing may be required is with subscriber managed QoS environments using a subscriber per SAP model. Usually, ingress multipoint traffic is minimal per subscriber and the extra multipoint queues for each subscriber reduces the overall subscriber density on the ingress forwarding plane. Multipoint shared queuing eliminates the multipoint queues sparing hardware queues for better subscriber density. Figure 2.3 demonstrates multipoint shared queuing.

One disadvantage of enabling multipoint shared queuing is that multipoint packets are no longer managed per service (although the unicast forwarding queues may provide limited benefit in this

area). Multipoint packets in a multipoint service (VPLS, IES and VPRN) use significant resources in the system, consuming ingress forwarding plane multicast bandwidth and egress replication bandwidth. Usually, the per service unicast forwarding queues are not rate limited to a degree that allows adequate management of multipoint packets traversing them when multipoint shared queuing is enabled. It is possible to minimize the amount of aggregate multipoint bandwidth by setting restrictions on the multipoint queue parameters in the QoS node's Shared Queue policy. Aggregate multipoint traffic can be managed per forwarding class for each of the three forwarding types (broadcast, multicast or unknown unicast – broadcast and unknown unicast are only used by VPLS).

A second disadvantage to multipoint shared queuing is the fact that multipoint traffic now consumes double the ingress forwarding plane bandwidth due to dual pass ingress processing.

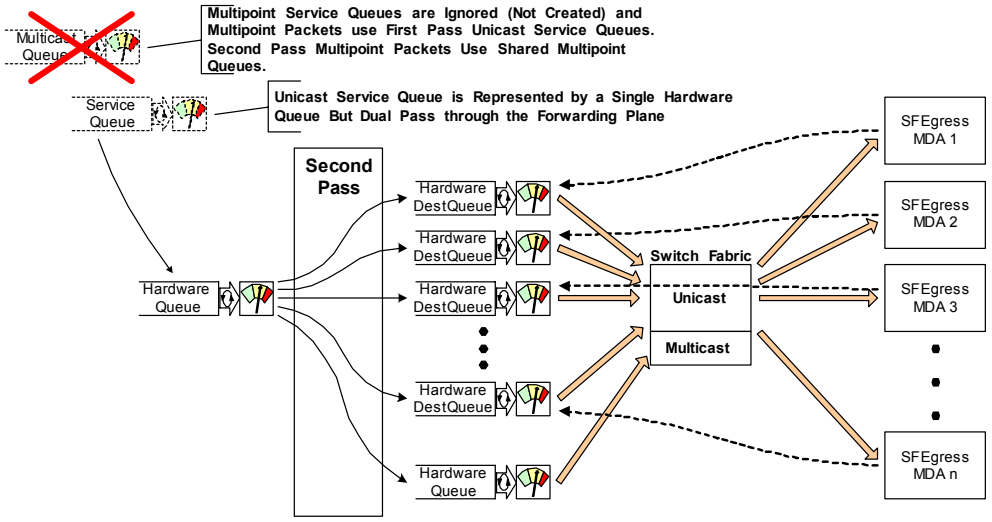

Figure 4: Multipoint Shared Queuing Using First Pass Unicast Queues

Alcatel-Lucent Service Model

In the SR-Series service model, the SR-Series service edge routers are deployed at the provider edge. Services are provisioned on SR-Series and transported across an IP and/or IP/MPLS provider core network in encapsulation tunnels created using Generic Router Encapsulation (GRE) or MPLS Label Switched Paths (LSPs).

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

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

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

Service Entities

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

- Customers (see page 36)
- Service Access Points (SAPs) (see page 36)
- Service Distribution Points (SDPs) (see page 41) (for distributed services only)

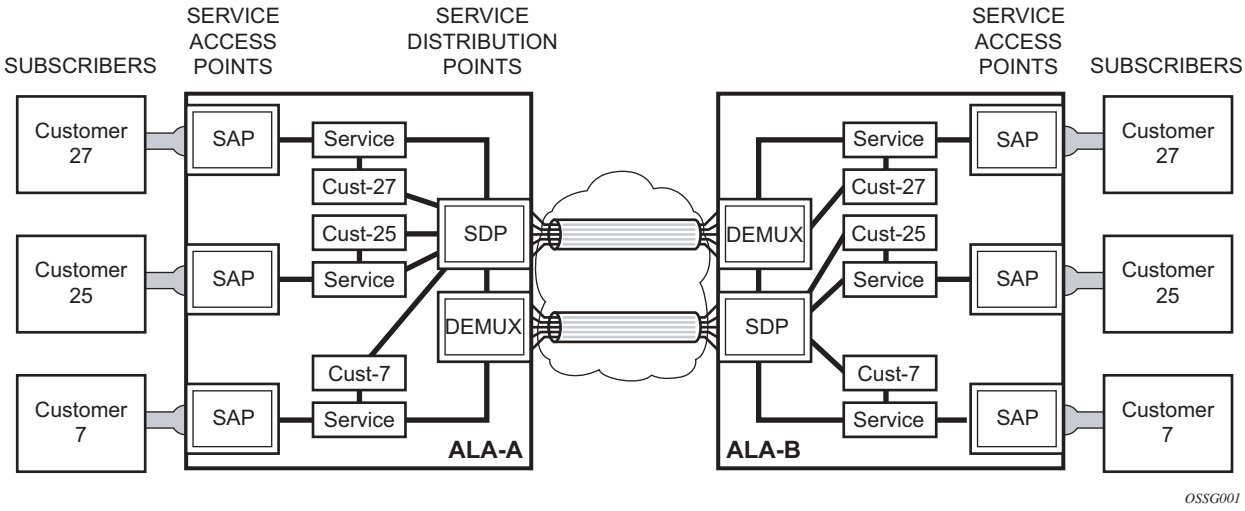

Figure 5: Service Entities

Customers

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

Service Access Points (SAPs)

Each subscriber service type is configured with at least one service access point (SAP). A SAP identifies the customer interface point for a service on an Alcatel-Lucent SR-Series router (Figure 6). The SAP configuration requires that slot, MDA, and port/channel information be specified. The slot, MDA, and port/channel parameters must be configured prior to provisioning a service (see the Cards, MDAs, and Ports sections of the *7750 SR OS Interface Configuration Guide*).

A SAP is a local entity to the 7750 SRand is uniquely identified by:

- The physical Ethernet port or SONET/SDH port or TDM channel
- The encapsulation type
- The encapsulation identifier (ID)

Depending on the encapsulation, a physical port or channel can have more than one SAP associated with it. SAPs can only be created on ports or channels designated as “access” in the physical port configuration. SAPs cannot be created on ports designated as core-facing “network” ports as these ports have a different set of features enabled in software.

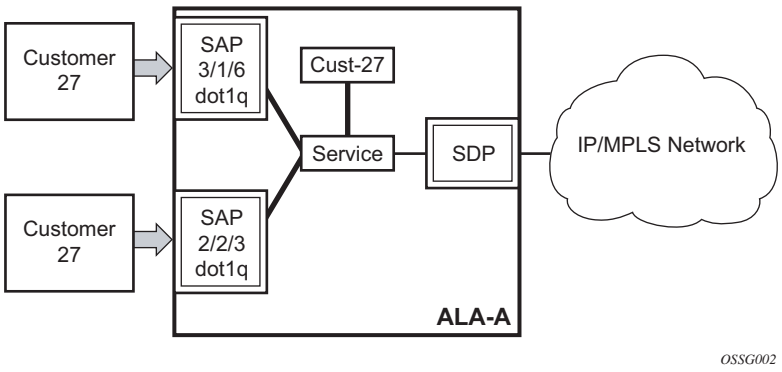

Figure 6: Service Access Point (SAP)

SAP Encapsulation Types and Identifiers

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

Ethernet Encapsulations

The following lists encapsulation service options on Ethernet ports:

- Null — Supports a single service on the port. For example, where a single customer with a single service customer edge (CE) device is attached to the port. The encapsulation ID is always 0 (zero).
- Dot1q — Supports multiple services for one customer or services for multiple customers (Figure 7). For example, the port is connected to a multi-tenant unit (MTU) device with multiple downstream customers. The encapsulation ID used to distinguish an individual service is the VLAN ID in the IEEE 802.1Q header.
- QinQ — The qinq encapsulation type adds a IEEE 802.1Q tag to the 802.1Q tagged packets entering the network to expand the VLAN space by tagging tagged packets, producing a double tagged frame.

SONET/SDH Encapsulations

There are several encapsulation service options on SONET/SDH channels:

- Internet Protocol Control Protocol (IPCP) — Supports a single IP service on a SONET/SDH port or a single service per channel (if the interface is channelized). This is typically used for router interconnection using point-to-point protocol (PPP).
- Bridging Control Protocol (BCP-null) — Supports a single service on the SONET/SDH port or a single service per channel (if the interface is channelized). This is used for bridging a single service between two devices using PPP over SONET/SDH. The encapsulation ID is always 0 (zero).
- Bridging Control Protocol (BCP-dot1q) — Supports multiple services on the SONET/SDH port/channel. This encapsulation type is used for bridging multiple services between two devices using PPP over SONET/SDH. The encapsulation ID used to distinguish services is the VLAN ID in the IEEE 802.1Q header in the BCP-encapsulated frame.
- ATM — ATM, ATM-FR, ATM SAP-bridge encapsulation termination Epipe and VPLS.

- Frame Relay — Supports the switched data link layer protocol that handles multiple virtual circuits.

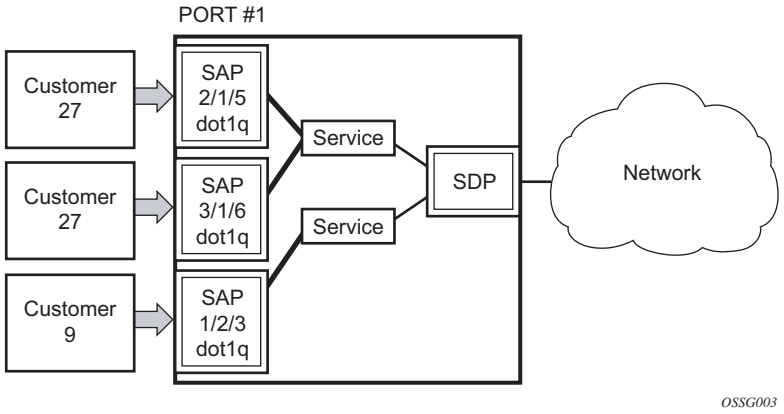

Figure 7: Multiple SAPs on a Single Port/Channel

Default SAP on a Dot1q Port

This feature introduces default SAP functionality on dot1q-encapsulated ports. This is similar to the functionality provided by Q1* SAP on q-in-q encapsulated ports, meaning that on dot1q-encapsulated ports where a default SAP is configured, all packets with q-tags not matching any explicitly defined SAPs will be assigned to this SAP.

One of the applications where this feature can be applicable is an access connection of a customer who uses the whole port to access Layer 2 services. The internal VLAN tags are transparent to the service provider. This can be provided by a null encapsulated port. A dedicated VLAN (not used by the customer) can be used to provide CPE management.

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

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

- This type of SAP is supported only on VPLS and EPIPE services and cannot be created in IES and VPRN services as it cannot preserve VLAN tag markings.
- For VPLS SAPs with STP enabled, STP listens to untagged and null-tagged BPDUs only. All other tagged BPDUs are forwarded like other customer packets. This is the same behavior as null-encapsulated ports.

- IGMP snooping is not supported on a default SAP. This would require remembering VLAN tags per hosts. By not allowing IGMP snooping of this SAP, all IGMP packets will be transparently forwarded.
- This type of SAP is mutually exclusive with a SAP defined by explicit null encapsulation (for example, 1/1/1:0). This avoids conflict as to which SAP untagged frames should be associated.

Services and SAP Encapsulations

Port Type	Encapsulation
Ethernet	Null
Ethernet	Dot1q
Ethernet	Q-in-Q
SONET/SDH	IPCP
SONET/SDH	BCP-null
SONET/SDH	BCP-dot1q
SONET/SDH	ATM
SONET/SDH	Frame Relay
SONET/SDH	Cisco HDLC

SAP Configuration Considerations

When configuring a SAP, consider the following:

- A SAP is a local entity and only locally unique to a given device. The same SAP ID value can be used on another SR-Series.
- There are no default SAPs. All SAPs in subscriber services must be created.
- The default administrative state for a SAP at creation time is administratively enabled.
- When a SAP is deleted, all configuration parameters for the SAP will also be deleted. For Internet Ethernet Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed.
- A SAP is owned by and associated with the service in which it is created in each SR-Series.
- A port/channel with a dot1q or BCP-dot1q encapsulation type means the traffic for the SAP is identified based on a specific IEEE 802.1Q VLAN ID value. The VLAN ID is stripped off at SAP ingress and the appropriate VLAN ID placed on at SAP egress. As a result, VLAN IDs only have local significance, so the VLAN IDs for the SAPs for a service need not be the same at each SAP.
- If a port/channel is administratively shutdown, all SAPs on that port/channel will be operationally out of service.
- A SAP cannot be deleted until it has been administratively disabled (shutdown).
- Each SAP can have one each of the following policies assigned:
  - → Ingress filter policy
  - → Egress filter policy
  - → Ingress QoS policy
  - → Egress QoS policy
  - → Accounting policy
  - → Ingress scheduler policy
  - → Egress scheduler policy

Service Distribution Points (SDPs)

A service distribution point (SDP) acts as a logical way to direct traffic from one SR-Series to another SR-Series through a uni-directional (one-way) service tunnel. The SDP terminates at the far-end SR-Series which directs packets to the correct service egress SAPs on that device. A distributed service consists of a configuration with at least one SAP on a local node, one SAP on a remote node, and an SDP binding the service to the service tunnel.

An SDP has the following characteristics:

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

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

SDP Binding

To configure a distributed service from ALA-A to ALA-B, the SDP ID (1) (Figure 8) must be specified in the service creation process in order to “bind” the service to the tunnel (the SDP). Otherwise, service traffic is not directed to a far-end point and the far-end SR-Series device(s) cannot participate in the service (there is no service). To configure a distributed service from ALA-B to ALA-A, the SDP ID (5) must be specified.

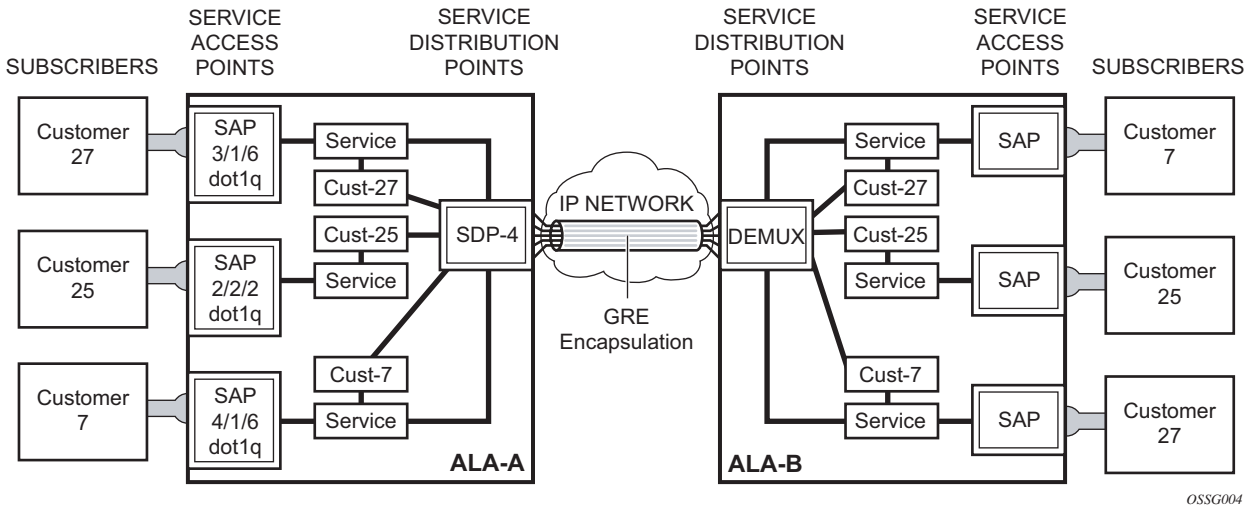

Figure 8: A GRE Service Distribution Point (SDP) pointing from ALA-A to ALA-B

Spoke and Mesh SDPs

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

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.

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

SDP Encapsulation Types

The Alcatel-Lucent service model uses encapsulation tunnels through the core to interconnect SR-Series service edge routers. An SDP is a logical way of referencing the entrance to an encapsulation tunnel.

The following encapsulation types are supported:

- L2 within Generic Routing Encapsulation (GRE)
- L2 within RSVP signaled, loose hop non-reserved MPLS LSP
- L2 within RSVP signaled, strict hop non-reserved MPLS LSP
- L2 within RSVP-TE signaled, bandwidth reserved MPLS LSP
- L2 with LDP signaled

GRE

GRE encapsulated tunnels have very low overhead and are best used for Best-Effort class of service. Packets within the GRE tunnel follow the IGP (Interior Gateway Protocol) shortest path from edge to edge. If a failure occurs within the service core network, the tunnel will only converge as fast as the IGP itself. If ECMP (Equal Cost Multi-Path) routing is used in the core, many loss-of-service failures can be minimized to sub-second timeframes.

MPLS

Multi-Protocol Label Switching (MPLS) encapsulation has the following characteristics:

- LSPs (label switched paths) are used through the network, for example, primary, secondary, loose hop, etc. These paths define how traffic traverses the network from point A to B. If a path is down, depending on the configuration parameters, another path is substituted.

  Paths can be manually defined or a constraint-based routing protocol (e.g., OSPF-TE or CSPF) can be used to determine the best path with specific constraints.
- An MPLS SR-Series router supports both signaled and non-signaled LSPs through the network.
- Non-signaled paths are defined at each hop through the network.
- Signaled paths are communicated via protocol from end to end using Resource ReserVation Protocol (RSVP).

  Because services are carried in encapsulation tunnels and an SDP is an entrance to the tunnel, an SDP has an implicit Maximum Transmission Unit (MTU) value. The MTU for

the service tunnel can affect and interact with the MTU supported on the physical port where the SAP is defined.

SDP Keepalives

SDP keepalives is a way of actively monitoring the SDP operational state using periodic Alcatel-Lucent SDP Ping Echo Request and Echo Reply messages. Alcatel-Lucent SDP Ping is a part of Alcatel-Lucent's suite of Service Diagnostics built on an Alcatel-Lucent service-level OA&M protocol. When SDP Ping is used in the SDP keepalive application, the SDP Echo Request and Echo Reply messages are a mechanism for exchanging far-end SDP status.

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

- AdminUp/AdminDown State
- Hello Time
- Message Length
- Max Drop Count
- Hold Down Time

SDP keepalive Echo Request messages are only sent when the SDP is completely configured and administratively up and SDP keepalives is administratively up. If the SDP is administratively down, keepalives for the SDP are disabled.

SDP keepalive Echo Request messages are sent out periodically based on the configured Hello Time. An optional Message Length for the Echo Request can be configured. If Max Drop Count Echo Request messages do not receive an Echo Reply, the SDP will immediately be brought operationally down.

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

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

For information about configuring keepalive parameters, refer to Configuring an SDP on page 68.

Class-Based Forwarding

Application of Class-Based Forwarding over RSVP LSPs

Class based forwarding over RSVP LSPs allows a service packet to be forwarded over a specific RSVP LSP, part of an SDP, based on its ingress determined forwarding class. The LSP selected depends on the operational status and load-balancing algorithms used for ECMP and LAG spraying.

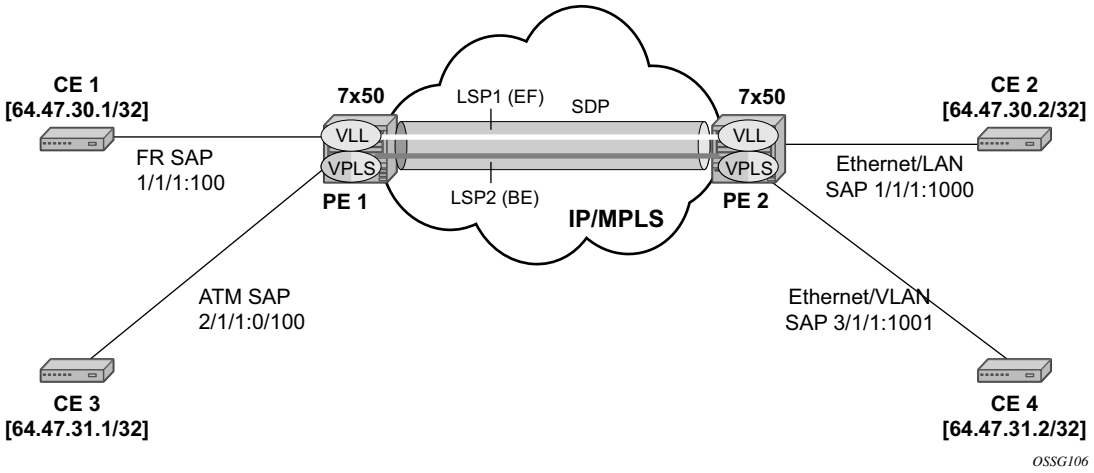

Figure 9: Class-Based Forwarding over SDP LSPs

Figure 9 illustrates the use of class-based forwarding to direct packets of a service to specific RSVP or static LSPs that are part of the same SDP based on the packets' forwarding class. The forwarding class of the packet is the one assigned to the packet as a result of applying the ingress QoS policy to the service SAP. The VLL service packets are all classified into the "`ef`" forwarding class and those that are destined to PE2 are forwarded over LSP1. Multicast and broadcast are classified into the "`be`" class and are forwarded over LSP2.

This feature allows service providers to dedicate specific LSPs with a determined level of traffic engineering and protection to select service packets. For example, packets of a VoIP service are assigned the "`ef`" class to expedite their forwarding but are also sent over carefully traffic-engineered and FRR-protected LSP paths across the service provider network.

Operation of Class-Based Forwarding over SDP LSPs

The 7750 SR class-based forwarding feature applies to a set of LSPs that are part of the same SDP. Each LSP must be configured as part of an SDP specifying the forwarding classes it will support. A forwarding class can only be assigned to one LSP in a given SDP, meaning that only one LSP within an SDP will support a given class of service. However, multiple classes of services can be assigned to an LSP. Both RSVP and static LSPs are allowed. All subclasses will be assigned to the same LSP as the parent forwarding class.

When a service packet is received at an ingress SAP, it is classified into one of the eight 7750 SR forwarding classes. If the packet will leave the SR on an SDP that is configured for class-based forwarding, the outgoing LSP will be selected based on the packet's forwarding class. Each SDP has a default LSP. The default LSP is used to forward a received packet that was classified at the ingress SAP into a forwarding class for which the SDP does not have an explicitly-configured LSP association. It is also used to forward a received packet if the LSP supporting its forwarding class is down. Note that the SDP goes down if the default LSP is down.

Class-based forwarding can be applied to all services supported by the 7750 SR. For VPLS services, explicit FC-to-LSP mappings are used for known unicast packets. Multicast and broadcast packets use the default LSP. There is a per-SDP user configuration that optionally overrides this behavior to specify an LSP to be used for multicast/broadcast packets.

VLL service packets are forwarded based on their forwarding class only if shared queuing is enabled on the ingress SAP. Otherwise, the class-forwarding command will have no effect and the VLL packets will be forwarded by hashing the service ID as per the hash algorithm used for LAG and ECMP.

Multi-Service Sites

A customer site can be designated a multi-service site where a single scheduler policy is applied to all SAPs associated with the site while retaining per-service and per-forwarding class shaping and policing. The SAPs associated with the multi-service site can be on a single port or on a single slot. The SAPs in a multi-service site cannot span slots.

Multi-service sites are anchor points to create an ingress and egress virtual scheduler hierarchy. When a site is created, it must be assigned to a chassis slot or port with the exception of the 7750 SR-1 in which the slot is set to 1. When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).

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.

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.

Each customer site must have a unique name within the context of the customer. Modifications made to an existing site immediately affect all SAPs associated with the site. Changing a scheduler policy association can 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.

Service Creation Process Overview

Figure 10 displays the overall process to provision core and subscriber services.

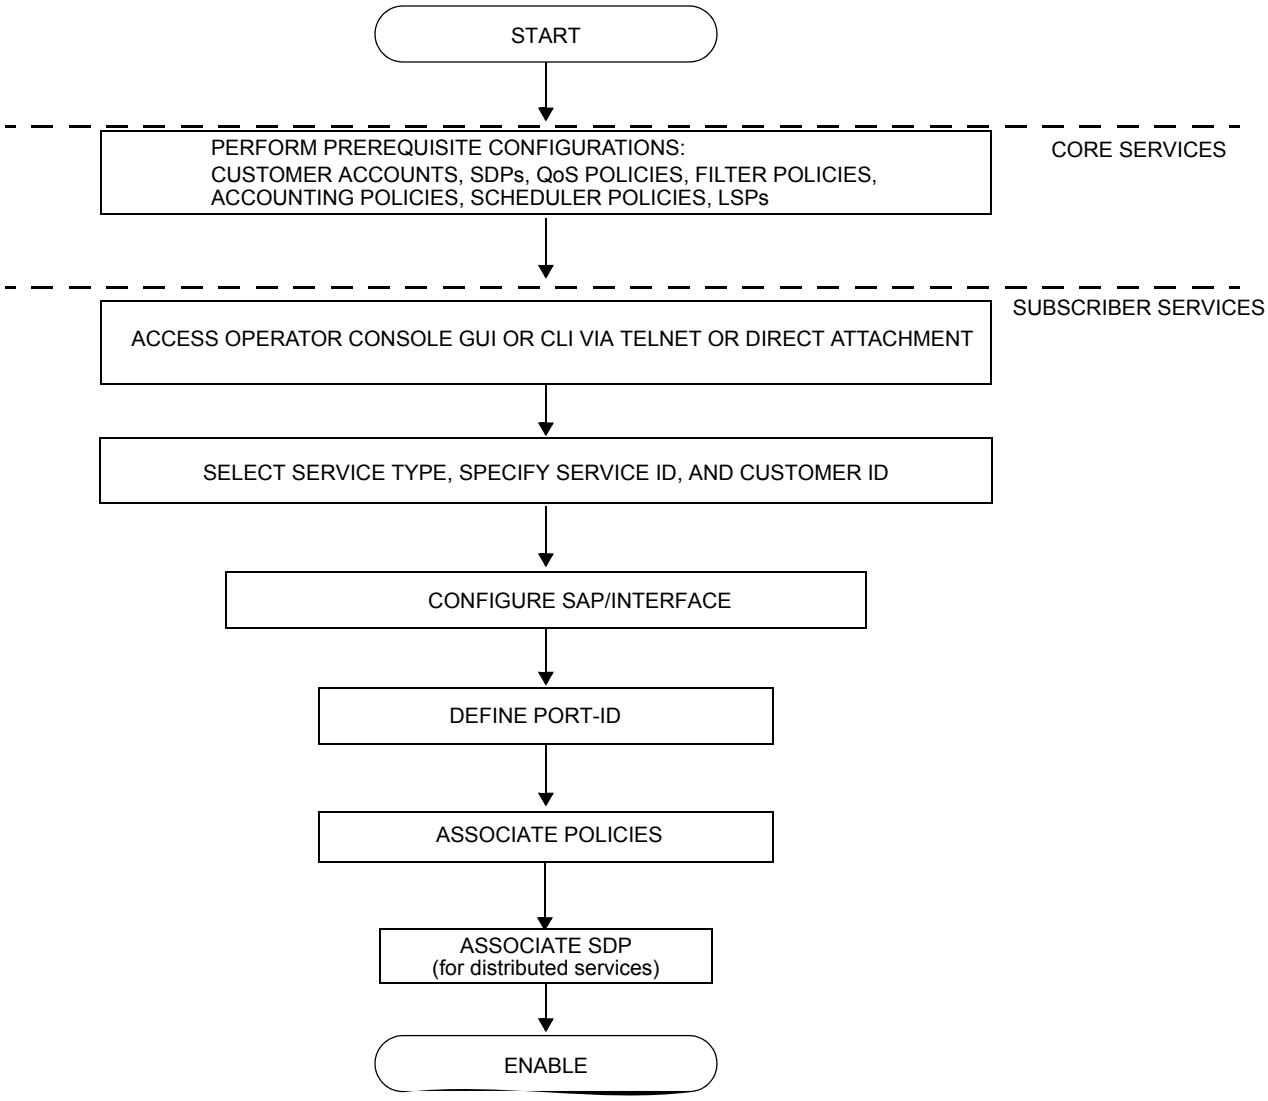

Figure 10: Service Creation and Implementation Flow

Service Components

Figure 11 displays the basic components of a subscriber service.

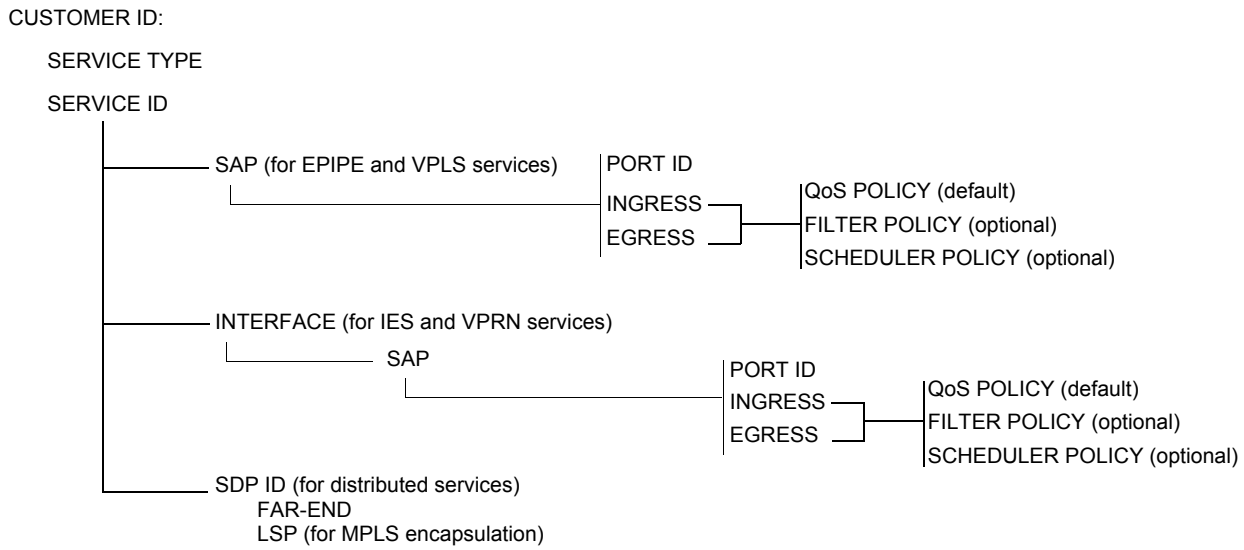

Figure 11: Subscriber Service Components

- Customer ID — This is the subscriber ID used to associate information with a particular customer.
- Service type — The service type specifies the connectivity type such as Epipe, VPLS, or IES.
- Service ID — Each service is uniquely identified by a Service ID within a service area.
- SAP — The subscriber-side service entry (access) and exit point for a service on an Alcatel-Lucent 7750 SR router.
- Interface — Configure a logical IP routing interface for Internet Ethernet Service (IES), Virtual Private LAN Service (VPLS), and Virtual Private Routed Network (VPRN) services.
- Port ID — The physical port identifier portion of the SAP definition.
- QoS policy — The QoS policy associated with an ingress or egress SAP or IP interface. QoS policy ID 1 is the default.
- Filter policy — Optional. The IP or MAC filter policy associated with an ingress or egress SAP or IP interface.
- Scheduler policy — Optional. A scheduler policy defines the hierarchy and operating parameters for virtual schedulers. This implementation of hierarchical scheduling is a method that defines a bounded operation for a group of queues. One or more queues are

mapped to a given scheduler with strict and weighted metrics controlling access to the scheduler. The scheduler has an optional prescribed maximum operating rate that limits the aggregate rate of the child queues. This scheduler may then feed into another virtual scheduler in a higher tier.

- SDP — An SDP is a logical mechanism that ties a far-end SR-Series to a specific service without having to specifically define the far-end SAPs. Each SDP, identified by a local SDP ID, represents a method for reaching a far-end SR-Series service edge router.

Deploying and Provisioning Services

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

Phase 1: Core Network Construction

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

- Build the IP or IP/MPLS core network.
- Configure routing protocols.
- Configure MPLS LSPs (if MPLS is used).
- Construct the core SDP service tunnel mesh for the services.

Phase 2: Service Administration

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

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

Phase 3: Service Provisioning

- Provision customer account information.
- If necessary, build any customer-specific QoS, filter or accounting policies.
- Provision the customer services on the SR-Series service edge routers by defining SAPs, binding policies to the SAPs, and then binding the service to appropriate SDPs as necessary. Refer to Configuring Customers on page 64 and Configuring an SDP on page 68.

Configuration Notes

This section describes service configuration caveats.

General

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

Core tasks include the following:

- Create customer accounts
- Create template QoS, filter, scheduler, and accounting policies
- Create LSPs
- Create SDPs

Subscriber services tasks include the following:

- Create Epipe, Ipipe, Apipe, Fpipe, VPLS, IES, or VPRN services
- Configure interfaces (where required) and SAPs
- Bind SDPs
- Create exclusive QoS and filter policies

Reference Sources

For information on standards and supported MIBs, refer to Standards and Protocol Support on page 1471.

Configuring Global Service Entities with CLI

This section provides information to create subscriber (customer) accounts and configure Service Distribution Points (SDPs) using the command line interface.

Topics include:

Service Model Entities

The Alcatel-Lucent service model uses logical entities to construct a service. The service model contains four main entities to configure a service.

- Subscribers on page 64
- Service Distribution Points (SDPs) on page 68
- Services:
  - → Ethernet Pipe (Epipe) Services on page 110
  - → ATM VLL (Apipe) Services on page 113
  - → Frame Relay VLL (Fpipe) Services on page 118
  - → IP Interworking VLL (Ipipe) Services on page 122
  - → VPLS on page 361
  - → IES on page 681
  - → VPRN on page 879
- Service Access Points (SAPs)
  - → Ethernet Pipe (Epipe) Services on page 110
  - → Apipe SAP on page 181
  - → Fpipe SAP on page 186
  - → VPLS SAP on page 386
  - → IES SAP on page 695
  - → VPRN Interface SAP on page 917

Service CLI Command Structure

Figure 12 displays the services CLI command structure. The service configuration commands are located under the **config>service** context.

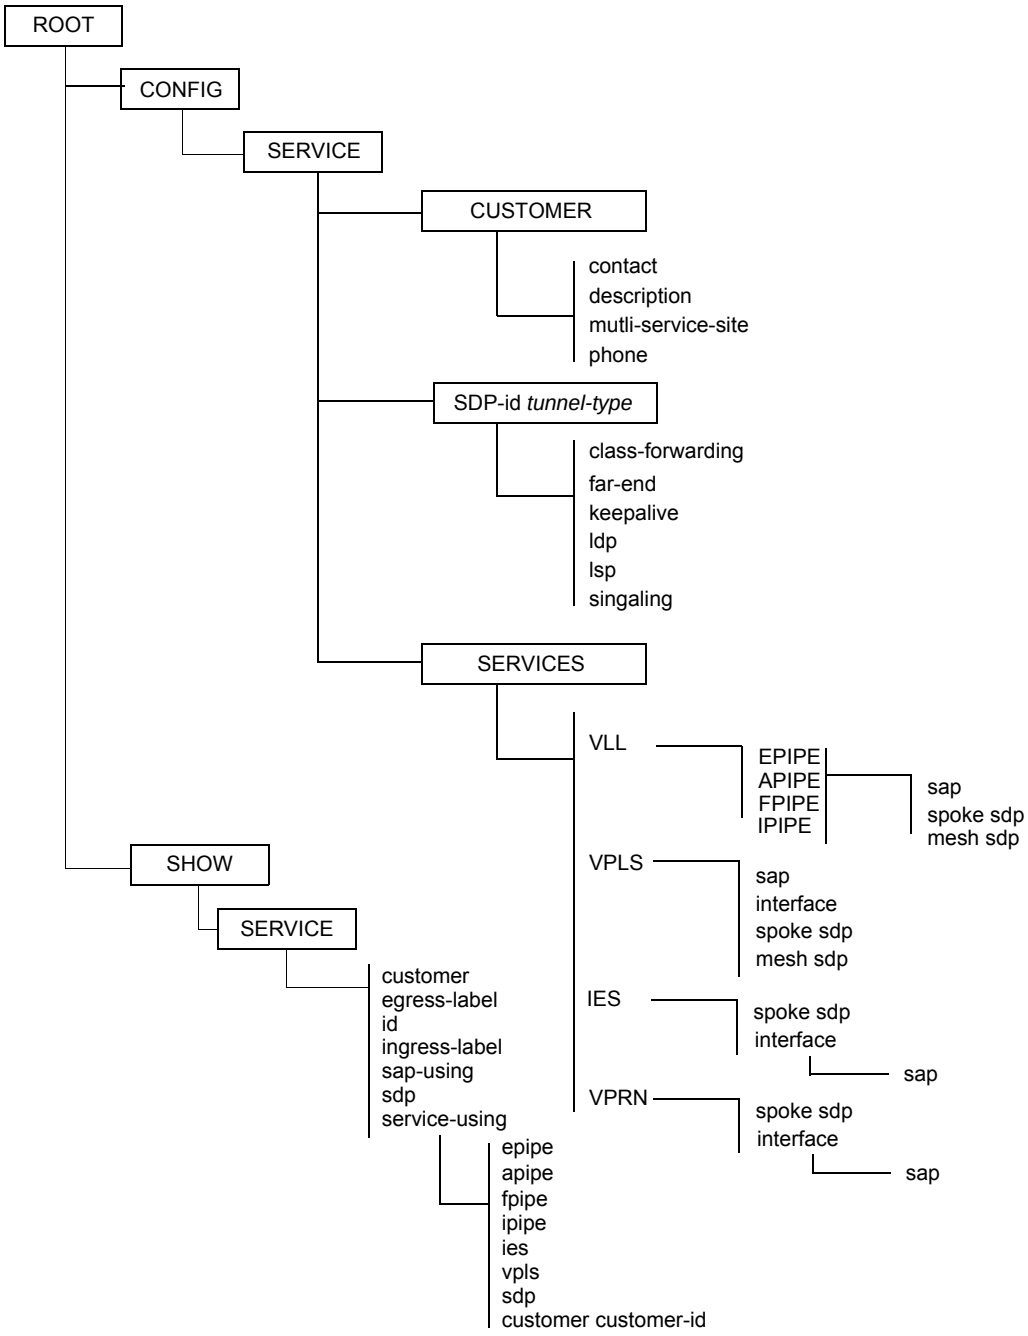

Figure 1: Global Service CLI Configuration Example

Figure 13 displays the CLI command structure for core and subscriber tasks which should be performed prior to provisioning a subscriber service.

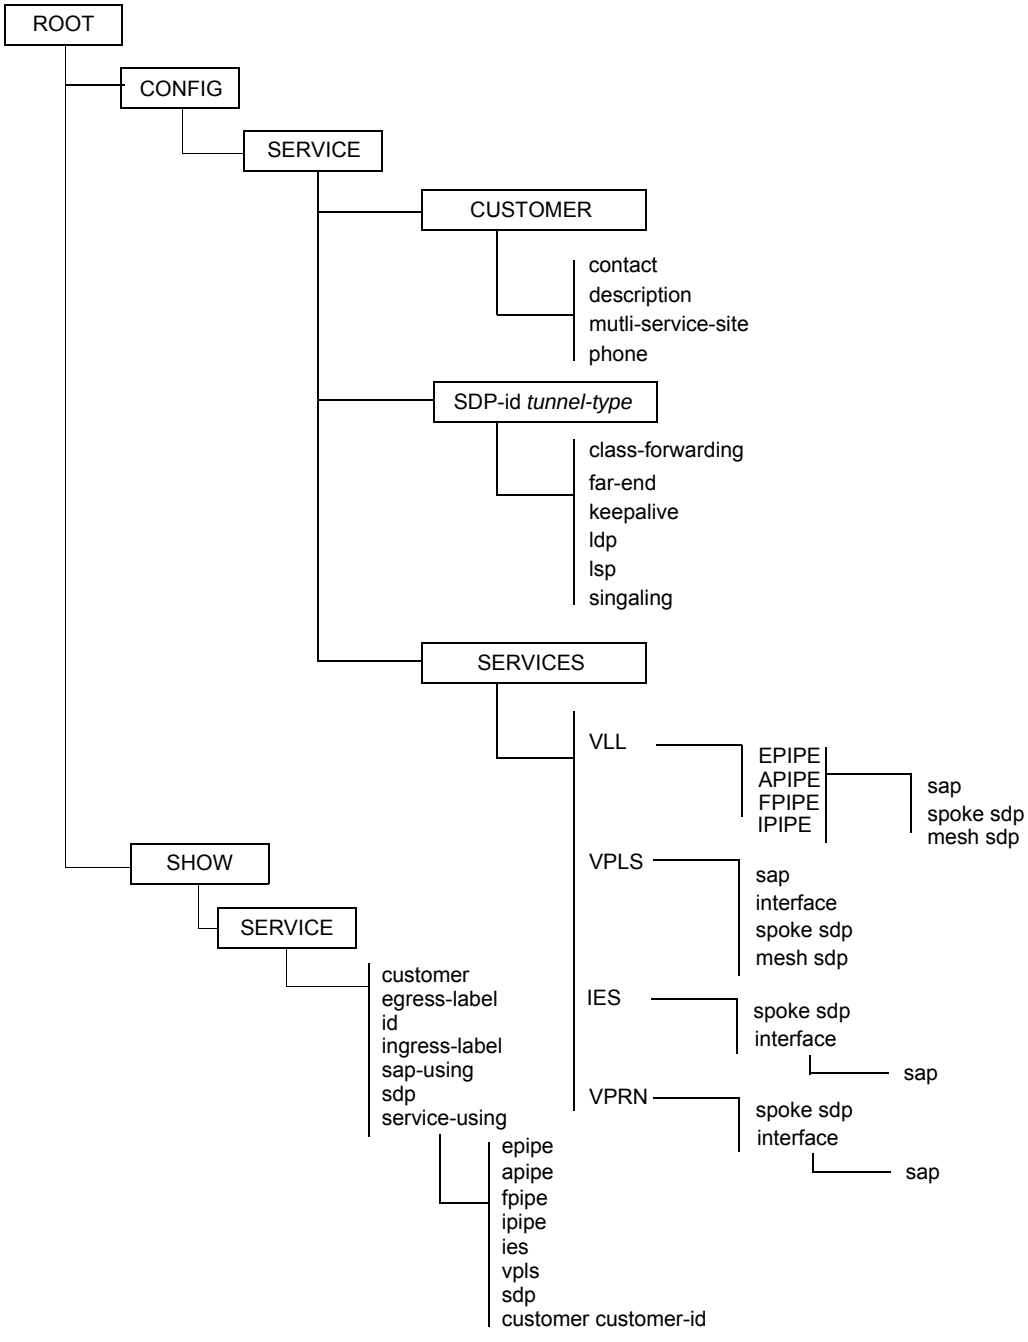

Figure 2: Core and Subscriber Tasks Configuration Example

List of Commands

Table 2 lists all the configuration commands to configure subscriber accounts and SDPs, indicating the configuration level at which each command is implemented with a short command description.

The services command list is organized in the following task-oriented manner:

- Configuring a customer account
- Configuring an SDP
  - → Configuring SDP class forwarding parameters
  - → Configure SDP Keepalive parameters
-

Table 1: CLI Commands to Configure Service Parameters

Command	Description	Page
Configuring a customer account		
`config>service`		64
`customer`	Creates a customer ID and customer context that is used to associate information with a particular customer.	81
`contact`	Creates a customer ID and customer context that is used to associate information with a particular customer.	81
`description`	Creates a text description stored in the configuration file for the customer.	79
`multi-service-site`	Creates a new customer site or edits an existing customer site. A customer site is an anchor point to create an ingress and egress virtual scheduler hierarchy. Multi-service customer sites create a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).	82
`assignment`	Specifies the port for the multi-service.	83
`description`	Creates a text description stored in the configuration file for a configuration context for the multi-service-site.	79
`egress`	Specifies egress parameters and policies.	84
`agg-rate-limit`	Defines a maximum total rate for all egress queues on a service SAP or multi-service site.	85
`ingress`	Specifies ingress parameters and policies.	84
`tod-suite`	Configures a time-of-day suite for this multi-service site.	89

Table 1: CLI Commands to Configure Service Parameters (Continued)

Command	Description	Page
`scheduler-override`	Specifies the set of attributes whose values have been overridden via management on this virtual scheduler.	85

Table 1: CLI Commands to Configure Service Parameters (Continued)

Command	Description	Page
scheduler	Defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler	86
rate	Defines the maximum bandwidth that the scheduler can offer its child queues or schedulers.	87
scheduler-policy	Applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues associated with this multi-service customer site.	88
phone	Adds phone number information for a customer ID.	82
Configuring an SDP		68
config>service# sdp *sdp-id* [gre|mpls]		90
gre	Specifies that the SDP will use GRE to reach the far-end router.	90
mpls	Specifies the SDP will use MPLS encapsulation and one or more LSP tunnels to reach the far-end router.	91
adv-mtu-override	Overrides the advertised VC MTU.	91
far-end	Configures the system IP address of the far-end destination router for the SDP that is terminating services.	93
ldp	Enables LDP-signaled LSP's on MPLS-encapsulated SDPs.	94
lsp	Creates associations between one or more label switched paths (LSPs) and an MPLS SDP.	94
path-mtu	Configures the Maximum Transmission Unit (MTU) in bytes that the SDP can transmit to the far-end router without packet dropping the SDP-type default path-mtu.	96
signaling	Enables the signaling protocol (targeted LDP) to obtain the ingress and egress labels in frames transmitted and received on the SDP.	95
vlan-vc-etype	Specifies the VLAN VC EtherType.	96
no shutdown	Administratively enables the SDP.	79
Configuring SDP class forwarding parameters		
config>service>sdp		
class-forwarding	Enables the forwarding of a service packet over the SDP based on the class of service of the packet.	92
fc	Creates an explicit association between a forwarding class and an LSP	92
multicast-lsp	Specifies the RSVP or static LSP in this SDP to use to forward VPLS multicast and broadcast packets.	93
no shutdown	Administratively enables class forwarding.	79

Table 1: CLI Commands to Configure Service Parameters (Continued)

Command	Description	Page
Configure SDP Keepalive parameters		
`config>service>sdp# lsp` *`lsp-name`*		
`keep-alive`	Configures SDP connectivity monitoring keepalive messages for the SDP ID.	97
`hello-time`	Configures the time period between SDP keepalive messages on the SDP-ID for the SDP connectivity monitoring messages.	98
`hold-down-time`	Configures the minimum time period the SDP will remain in the operationally down state in response to SDP keepalive monitoring.	98
`max-drop-count`	Configures the number of consecutive SDP keepalive failed request attempts or remote replies that can be missed after which the SDP is operationally downed.	98
`message-length`	Configures the size of SDP monitoring keepalive request messages.	99
`no shutdown`	Administratively enables the keepalive messages.	79

Basic Configuration

The most basic service configuration must have the following:

- A customer ID
- A service type
- A service ID
- A SAP identifying a port and encapsulation value
- An interface (where required) identifying an IP address, IP subnet, and broadcast address
- For distributed services: an associated SDP

The following example provides an Epipe service configuration displaying the SDP and Epipe service entities. SDP ID 2 was created with the far-end node 10.10.10.104. Epipe ID 6000 was created for customer ID 6 which uses the SDP ID 2.

```
A:ALA-B>config>service# info detail
#------------------------------------------
...
        sdp 2 gre create
            description "GRE-10.10.10.104"
            far-end 10.10.10.104
            signaling tldp
            no vlan-vc-etype
            keep-alive
            path-mtu 4462
            keep-alive
                shutdown
                hello-time 10
                hold-down-time 10
                max-drop-count 3
                timeout 5
                no message-length
            exit
            no shutdown
        exit
...
       epipe 6000 customer 6 vpn 6000 create
           service-mtu 1514
           sap 1/1/2:0 create
               no multi-service-site
               ingress
                   no scheduler-policy
                   qos 1
               exit
               egress
                   no scheduler-policy
                   qos 1
               exit
               no collect-stats
               no accounting-policy
               no shutdown
           exit
           spoke-sdp 2:6111 create
               ingress
```

```
                    no vc-label
                    no filter
                exit
                egress
                    no vc-label
                    no filter
                exit
                no shutdown
            exit
            no shutdown
        exit
...
#------------------------------------------
A:ALA-B>config>service#
```

Common Configuration Tasks

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

Configuring Customers

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

- Description
- Contact name
- Telephone number
- Multi-service site

Use the following CLI syntax to create and input customer information:

CLI Syntax:
```
config>service# customer customer-id create
    contact contact-information
    description description-string
    multi-service-site customer-site-name [create]
        assignment {port port-id | card slot}
        description description-string
        egress
            scheduler-override
                scheduler scheduler-name
                    rate pir-rate [cir cir-rate]
            scheduler-policy scheduler-policy-name
        ingress
            scheduler-override
                scheduler scheduler-name
                    rate pir-rate [cir cir-rate]
            scheduler-policy scheduler-policy-name
    phone phone-number
```

The following displays the configuration command usage to create a customer account:

Example:
```
config>service# customer 5 create
config>service>cust$ description "Alcatel Customer"
config>service>cust# contact "Technical Support"
config>service>cust# phone "650 555-5100"
config>service>cust# exit
```

The following displays the customer account configuration.

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

Configuring Multi-Service-Sites

Multi-service sites create a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs). The ingress and egress scheduler-policy commands on the SAP are mutually exclusive with the SAP `multi-service-site` command. The multi-service customer site association must be removed from the SAP before local scheduler polices may be applied.

After a multi-service site is created, it must be assigned to a chassis slot or port. Note that the 7750 SR-1 model multi-service site `assignment` configuration defaults to slot 1.

Use the following CLI syntax to configure customer multi-service sites.

CLI Syntax:

```
config>service> customer customer-id
    multi-service-site customer-site-name
       assignment {port port-id | card slot}
       description description-string
       egress
          agg-rate-limit agg-rate
          scheduler-policy scheduler-policy-name
       ingress
          scheduler-policy scheduler-policy-name
       tod-suite tod-suite-name
```

Example:

```
config>service# customer 5 create
     config>service>cust$ multi-service-site "WestCoast" create
     config>service>cust>multi-service-site$ assignment card 3
     config>service>cust>multi-service-site# egress scheduler-policy
SLA1
     config>service>cust>multi-service-site# exit
     config>service>cust# multi-service-site "EastCoast" create
     config>service>cust>multi-service-site$ assignment card 4
     config>service>cust>multi-service-site# ingress
     config>service>cust>multi-service-site>ingress$ scheduler-policy
alpha1
     config>service>cust>multi-service-site>ingress# exit
     config>service>cust>multi-service-site# exit
     config>service>cust# exit
```

The following displays a customer account configuration.

```
A:ALA-12>config>service# info
----------------------------------------
..
       customer 5 create
           multi-service-site "EastCoast" create
               assignment card 4
               ingress
                   scheduler-policy "alpha1"
               exit
           exit
           multi-service-site "WestCoast" create
               assignment card 3
               egress
                   scheduler-policy "SLA1"
               exit
           exit
           description "Alcatel Customer"
           contact "Technical Support"
           phone "650 555-5100"
       exit
...
----------------------------------------
A:ALA-12>config>service#
```

Configuring an SDP

The most basic SDP must have the following:

- A locally unique SDP identification (ID) number.
- The system IP address of the originating and far-end 7750 SR routers.
- An SDP encapsulation type - either GRE or MPLS.

SDP Configuration Tasks

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

Consider the following SDP characteristics:

- SDPs can be created as either GRE or MPLS.
- Each distributed service must have an SDP defined for every remote 7750 SR-Series router to provide VLL, VPLS, and VPRN services.
- A distributed service must be bound to an SDP. By default, no SDP is associated with a service. Once an SDP is created, services can be associated to that SDP.
- An SDP is not specific or exclusive to any one service or any type of service. An SDP can have more than one service bound to it.
- The SDP IP address must be a 7750 SR-Series system IP address.
- In order to configure an MPLS SDP, LSPs must be configured first and then the LSP-to-SDP association must be explicitly created.
- In the SDP configuration, automatic ingress and egress labeling (targeted LDP) is enabled by default. Ingress and egress VC labels are signaled over a TLDP connection between two 7750 SR-Series routers.

  Note that if signaling is disabled for an SDP, then services using that SDP must configure ingress and egress vc-labels manually.

To configure a basic SDP, perform the following steps:

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

Configuring an SDP

Use the following CLI syntax to create an SDP and select an encapsulation type. If you do not specify GRE or MPLS, the default encapsulation type is GRE.

NOTE: When you specify the far-end ip address, you are creating the tunnel. In essence, you are creating the path from Point A to Point B. When you configure a distributed service, you must identify an SDP ID. Use the show service sdp command to display the qualifying SDPs.

When specifying MPLS SDP parameters, you can only specify an LSP or enable LDP. There cannot be 2 methods of transport in a single SDP. If an LSP name is specified, then RSVP is used for dynamic signaling within the LSP.

LSPs are configured in the `config>router>mpls` context. See the *7750 SR OS MPLS Guide* for configuration and command information.

Use the following CLI syntax to create a GRE or MPLS SDP:

CLI Syntax:
```
config>service>sdp sdp-id [gre | mpls] create
    adv-mtu-override
    description description-string
    far-end ip-addr
    keep-alive
       hello-time seconds
       hold-down-time seconds
       max-drop-count count
       message-length octets
       timeout timeout
       no shutdown
             ldp                    (only for MPLS SDPs)
             lsp lsp-name [lsp-name](only for MPLS SDPs)
    path-mtu octets
    signaling {off|tldp}
    no shutdown
```

The following example display the syntax used to create a GRE SDP, an LSP-signalled MPLS SDP, and an LDP-signalled MPLS SDP.

Example:
```
config>service# sdp 2 gre create
config>service>sdp# description "to-GRE-10.10.10.104"
config>service>sdp# far-end "10.10.10.104"
config>service>sdp# no shutdown
config>service>sdp# exit
config>service# sdp 8 mpls create
config>service>sdp# description "MPLS-10.10.10.104"
config>service>sdp# 10.10.10.104
config>service>sdp# lsp "to-104"
```

```
            config>service>sdp# no shutdown
            config>service>sdp# exit

            config>service# sdp 104 mpls create
            config>service>sdp# description "MPLS-10.10.10.94"
            config>service>sdp# 10.10.10.94
            config>service>sdp# ldp
            config>service>sdp# no shutdown
            config>service>sdp# exit
```

The following displays the SDP sample configurations.

```
A:ALA-12>config>service# info
----------------------------------------
...
        sdp 2  create
            description "GRE-10.10.10.104"
            far-end 10.10.10.104
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        sdp 8 mpls create
            description "MPLS-10.10.10.104"
            far-end 10.10.10.104
            lsp "to-104"
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        sdp 104 mpls create
            description "MPLS-10.10.10.94"
            far-end 10.10.10.94
            ldp
            keep-alive
                shutdown
            exit
            no shutdown
        exit
...
----------------------------------------
A:ALA-12>config>service#
```

Service Management Tasks

This section discusses the following service management tasks:

Modifying Customer Accounts

To access a specific customer account, you must specify the customer ID.
To display a list of customer IDs, use the `show service customer` command.
Enter the parameter (`description`, `contact`, `phone`) and then enter the new information.

CLI Syntax:

```
config>service# customer customer-id create
    [no] contact contact-information
    [no] description description-string
    [no] multi-service-site customer-site-name [create]
       assignment {port port-id | card slot}
       no assignment
       [no] description description-string
       egress
          scheduler-policy scheduler-policy-name
          no scheduler-policy
       ingress
          scheduler-policy scheduler-policy-name
          no scheduler-policy
       tod-suite tod-suite-name
    [no] phone phone-number
```

Example:

```
config>service#  customer 27 create
config>service>customer$ description "Western Division"
config>service>customer# contact "John Dough"
config>service>customer# no phone "(650) 237-5102"
```

Deleting Customers

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

CLI Syntax: `config>service# no customer customer-id`

Example:
```
config>service# epipe 5 customer 27 shutdown
config>service# epipe 9 customer 27 shutdown
config>service# no epipe 5
config>service# no epipe 9
config>service# no customer 27
```

Modifying SDPs

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

NOTE: Once created, you cannot modify the SDP encapsulation type.

CLI Syntax: `config>service# sdp` *`sdp-id`*

Example:
```
config>service# sdp 79
config>service>sdp# description "Path-to-107"
config>service>sdp# shutdown
config>service>sdp# far-end "10.10.10.107"
config>service>sdp# path-mtu 1503
config>service>sdp# no shutdown
```

Deleting SDPs

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

CLI Syntax: `config>service# no sdp 79`

Example:
```
config>service# epipe 5 spoke-sdp 79:5
config>service>epipe>sdp# shutdown
config>service>epipe>sdp# exit
```

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

Modifying LSPs

Although the CLI command descriptions are discussed in the MPLS chapter of this book, the syntax is repeated in this section as a quick reference.

To access a specific LSP, you must specify the LSP ID. To display a list of LSPs, use the `show router mpls lsp` command. Enter the parameter, such as `from`, `to`, `retry-timer`, `retry-limit`, `primary`, and `secondary`, and then enter the new information.

CLI Syntax: `config>router>mpls# lsp` *`lsp-name`*

Example:

```
config>router>mpls# lsp to-headquarters
config>router>mpls>lsp# fast-reroute hop-limit 35
config>router>mpls>lsp# primary to-104
config>router>mpls>lsp# bandwidth 50000
config>router>mpls>lsp# no shutdown
```

Deleting LSPs

The no form of the `lsp` command removes an LSP ID and all associated information. Before an LSP can be deleted, the LSP must be removed from all SDP associations. The SDP must be administratively disabled before deleting LSPs.

CLI Syntax:

```
config>router# mpls
  [no] lsp lsp-name
  shutdown
```

CLI Syntax:

```
config>service# sdp sdp-id
  [no] lsp lsp-name
```

Example:

```
config>service# sdp 79
config>service>sdp# no lsp 123
config>service>sdp# exit all
# config router
config>router# mpls
config>router>mpls# lsp 123
config>router>mpls>lsp# shutdown
config>router>mpls>lsp# exit
config>router>mpls# no lsp 123
```

Global Services Command Reference

Command Hierarchies

Customer Commands

```
config
    — service
        — [no] customer customer-id
            — contact contact-information
            — no contact
            — description description-string
            — no description
            — multi-service-site customer-site-name
            — no multi-service-site customer-site-name
                — assignment {port port-id | card slot-number}
                — no assignment
                — description  description-string
                — no description
                — egress
                    — agg-rate-limit agg-rate
                    — no agg-rate-limit
                    — [no] scheduler-override
                        — [no] scheduler scheduler-name
                            — rate  pir-rate [cir cir-rate]
                            — no rate
                    — scheduler-policy scheduler-policy-name
                    — no scheduler-policy
                — ingress
                    — [no] scheduler-override
                        — [no] scheduler scheduler-name
                            — rate  pir-rate [cir cir-rate]
                            — no rate
                    — scheduler-policy scheduler-policy-name
                    — no scheduler-policy
                — tod-suite tod-suite-name
                — no tod-suite
            — [no] phone phone-number
```

SDP Commands

```
config
    — service
        — sdp sdp-id [gre | mpls]
        — no sdp sdp-id
            — accounting-policy acct-policy-id
            — no accounting-policy
            — [no] adv-mtu-override
            — class-forwarding [default-lsp lsp-name]
            — no class-forwarding
                — fc {be | l2 | af | l1 | h2 | ef | h1 | nc}  lsp lsp-name
                — no fc  {be | l2 | af | l1 | h2 | ef | h1 | nc}
                — multicast-lsp lsp-name
                — no multicast-lsp
                — [no] shutdown
            — [no] collect-stats
            — description description-string
            — no description
            — far-end ip-addr
            — no far-end
            — keep-alive
                — hello-time seconds
                — no hello-time
                — hold-down-time seconds
                — no hold-down-time
                — max-drop-count count
                — no max-drop-count
                — message-length octets
                — no message-length
                — [no] shutdown
                — timeout timeout
                — no timeout
            — [no] ldp
            — [no] lsp lsp-name
            — metric metric
            — no metric
            — path-mtu octets
            — no path-mtu
            — signaling [off | tldp]
            — vlan-vc-etype 0x0600..0xffff
            — no vlan-vc-etype [x0600.0xffff]
```

SAP Commands

```
config
    — service
        — apipe
            — sap sap-id [create] [no-endpoint]
            — sap sap-id [create] endpoint endpoint-name
            — no sap sap-id
        — epipe
            — sap sap-id [create] [no-endpoint]
            — sap sap-id [create] endpoint endpoint-name
            — no sap sap-id
        — fpipe
            — sap sap-id [create] [no-endpoint]
            — sap sap-id [create] endpoint endpoint-name
            — no sap sap-id
        — ies
            — sap sap-id [create]
            — no sap sap-id
        — ipipe
            — sap sap-id [create] [no-endpoint]
            — sap sap-id [create] endpoint endpoint-name
            — no sap sap-id
        — vpls
            — sap sap-id [split-horizon-group group-name] [create]
            — no sap sap-id
        — vprn
            — sap sap-id [create]
            — no sap sap-id
```

Show Commands

show
- **service**
  - **customer** [*customer-id*] [**site** *customer-site-name*]
  - **sdp** [*sdp-id* | **far-end** *ip-addr*] [**detail** | **keep-alive-history**]
  - **sdp-using** [*sdp-id*[:*vc-id*] | **far-end** *ip-address*]
  - **service-using** [**epipe**] [**ies**] [**vpls**] [**vprn**] [**mirror**] [**apipe**] [**fpipe**] [**ipipe**] [**sdp** *sdp-id*] [**customer** *customer-id*]

Global Service Configuration Commands

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>service>sdp
config>service>sdp>class-forwarding
config>service>sdp>keep-alive
config>service>sdp>forwarding-class

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 **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.

SDP (global) — When an SDP is shutdown at the global service level, all bindings to that SDP are put into the out-of-service state and the SDP itself is put into the administratively and operationally down states. Packets that would normally be transmitted using this SDP binding will be discarded and counted as dropped packets.

SDP (service level) — Shutting down an SDP within a service only affects traffic on that service from entering or being received from the SDP. The SDP itself may still be operationally up for other services.

SDP Keepalives — Enables SDP connectivity monitoring keepalive messages for the SDP ID. Default state is disabled (shutdown) in which case the operational state of the SDP-ID is not affected by the keepalive message state.

description

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

Context config>service>customer
config>service>customer>multi-service-site
config>service>sdp

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

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

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

Default No description associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters 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.

Customer Management Commands

customer

Syntax **customer** *customer-id*
no customer *customer-id*

Context config>service

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

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

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

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

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

Parameters *customer-id* — Specifies the ID number to be associated with the customer, expressed as an integer.

Values 1 — 2147483647

contact

Syntax **contact** *contact-information*
no contact *contact-information*

Context config>service>customer

Description This command allows you to configure contact information for a customer.

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

Default No contact information is associated with the *customer-id*.

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

Parameters *sontact-information* — The customer contact information entered as an ASCII character string up to 80 characters in length. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. Any printable, seven bit ASCII characters may be used within the string.

phone

Syntax [**no**] **phone** *string*

Context config>service>customer *customer-id*

Description This command adds telephone number information for a customer ID.

Default No telephone number information is associated with a customer.

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

Parameters *string* — The customer phone number entered as an ASCII string string up to 80 characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. Any printable, seven bit ASCII characters may be used within the string.

multi-service-site

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

Context config>service>customer

Description This command creates a new customer site or edits an existing customer site with the *customer-site-name* parameter. A customer site is an anchor point to create an ingress and egress virtual scheduler hierarchy. When a site is created, it must be assigned to a chassis slot or port with the exception of the 7750 SR-1 in which the slot is set to 1. When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).

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.

assignment

Syntax **assignment** {**port** *port-id* | **card** *slot-number*}
no assignment

Context config>service>customer>multi-service-site

Description This command assigns a multi-service customer site to a specific chassis slot, port, or channel. This allows the system to allocate the resources necessary to create the virtual schedulers defined in the ingress and egress scheduler policies as they are specified. This also verifies that each SAP assigned to the site exists within the context of the proper customer ID and that the SAP was configured on the proper slot, port, or channel. The assignment must be given prior to any SAP associations with the site.

The **no** form of the command removes the port, channel, or slot assignment. If the customer site has not yet been assigned, the command has no effect and returns without any warnings or messages.

Default **None**

Parameters **port** *port-id* — The **port** keyword is used to assign the multi-service customer site to the port-id or port-id.channel-id given. When the multi-service customer site has been assigned to a specific port or channel, all SAPs associated with this customer site must be on a service owned by the customer and created on the defined port or channel. The defined port or channel must already have been pre-provisioned on the system but need not be installed when the customer site assignment is made.

Syntax: *port-id*[:encap-val]

Values

port-id	*slot*/*mda*/*port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	

	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200

card *slot-number* — The **card** keyword is used to assign the multi-service customer site to the slot-number given. When the multi-service customer site has been assigned to a specific slot in the chassis, all SAPs associated with this customer site must be on a service owned by the customer and created on the defined chassis slot.

The defined slot must already have been pre-provisioned on the system but need not be installed when the customer site assignment is made.

Values Any pre-provisioned slot number for the chassis type that allows SAP creation
slot-number 1 — 10

ingress

Syntax ingress

Context config>service>customer>multi-service-site

Description This command enables the context to configure the ingress node associate an existing scheduler policy name with the customer site. The ingress node is an entity to associate commands that complement the association.

egress

Syntax egress

Context config>service>customer>multi-service-site

Description This command enables the context to configure the egress node associate an existing scheduler policy name with the customer site. The egress node is an entity to associate commands that complement the association.

agg-rate-limit

Syntax **agg-rate-limit** *agg-rate*
no agg-rate-limit

Context config>service>customer>multi-service-site>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, at 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

scheduler-override

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

Context config>service>customer>multi-service-site>ingress
config>service>customer>multi-service-site>egress

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.

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 queues 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**

scheduler

Syntax [**no**] **scheduler** *scheduler-name*

Context config>service>customer>multi-service-site>ingress>sched-override
config>service>customer>multi-service-site>egress>sched-override

Description This command can be used to override specific attributes of the specified scheduler name.

A scheduler defines 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>customer>multi-service-site>ingress>sched-override>scheduler
config>service>customer>multi-service-site>egress>sched-override>scheduler

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 scheduler's amount of bandwidth to be considered during the parent schedulers '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 queues 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**

scheduler-policy

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

Context config>service>customer>multi-service-site>ingress
config>service>customer>multi-service-site>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.

tod-suite

Syntax **tod-suite** *tod-suite-name*
no tod-suite

Context config>service>cust>multi-service-site

Description This command applies a time-based policy (filter or QoS policy) to the multiservice site. The suite name must already exist in the **config>cron** context.

Default **no tod-suite**

Parameters *tod-suite-name* — Specifies collection of policies (ACLs, QoS) including time-ranges. Only the scheduler-policy part of the tod-suite is taken into account. The suite can be applied to more than one multi-service-site.

SDP Commands

sdp

Syntax **sdp** *sdp-id* [{**gre** | **mpls**}]
no sdp *sdp-id*

Context config>service

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

An SDP is a logical mechanism that ties a far-end 7750 SR to a particular service without having to specifically define far end SAPs. Each SDP represents a method to reach a 7750 SR router.

One method is IP Generic Router Encapsulation (GRE) which has no state in the core of the network. GRE does not specify a specific path to the 7750 SR. A GRE-based SDP uses the underlying IGP routing table to find the best next hop to the far end 7750 SR.

The other method is Multi-Protocol Label Switching (MPLS) encapsulation. A 7750 SR supports both signaled and non-signaled Label Switched Paths (LSPs) through the network. Non-signaled paths are defined at each hop through the network. Signaled paths are communicated via protocol from end to end using Resource ReserVation Protocol (RSVP). Paths may be manually defined or a constraint-based routing protocol (i.e., OSPF-TE or CSPF) can be used to determine the best path with specific constraints.

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

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

If *sdp-id* exists, the current CLI context is changed to that SDP for editing and modification. For editing an existing SDP, neither the **gre** nor **mpls** keyword is specified. If a keyword is specified for an existing *sdp-id*, an error is generated and the context of the CLI will not be changed to the specified *sdp-id*.

The **no** form of this command deletes the specified SDP. Before an SDP can be deleted, it must be administratively down (shutdown) and not bound to any services. If the specified SDP is bound to a service, the **no sdp** command will fail generating an error message specifying the first bound service found during the deletion process. If the specified *sdp-id* does not exist an error will be generated.

Default **none**

Parameters *sdp-id* — The SDP identifier.

Values 1 — 17407

gre — Specifies the SDP will use GRE to reach the far-end 7750 SR. Only one GRE SDP can be created to a given destination 7750 SR. Multiple GRE SDPs to a single destination 7750 SR serve no purpose as the path taken to reach the far end 7750 SR is determined by the IGP which will be the same for all SDPs to a given destination and there is no bandwidth reservation in GRE tunnels.

mpls — Specifies the SDP will use MPLS encapsulation and one or more LSP tunnels to reach the far-end 7750 SR. Multiple MPLS SDPs may be created to a given destination 7750 SR. Multiple MPLS SDPs to a single destination 7750 SR are helpful when they use divergent paths.

accounting-policy

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

Context config>service>sdp

Description This command creates the accounting policy context that can be applied to an SDP.

An accounting policy must be defined before it can be associated with a SDP.
If the *policy-id* does not exist, an error message is generated.

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

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

Default Default accounting policy.

Parameters *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 — 99

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>sdp

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

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default **collect-stats**

adv-mtu-override

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

Context config>service>sdp

Description This command overrides the advertised VC-type MTU. When enabled, the SR-Series signals a VC MTU equal to the service MTU (includes the L2 header). Under normal operations it will advertise the service MTU minus the L2 header. In the receive direction, it will accept either one.

The **no** form of the command disables the VC-type MTU override.

Default **no adv-mtu-override**

class-forwarding

Syntax **class-forwarding** [**default-lsp** *lsp-name*]
no class-forwarding

Context config>service>sdp

Description This command enables the forwarding of a service packet over the SDP based on the class of service of the packet. Specifically, the packet is forwarded on the RSVP LSP or static LSP whose forwarding class matches that of the packet. The user maps the system forwarding classes to LSPs using the **config>service>sdp>class-forwarding>fc** command. If there is no LSP that matches the packet's forwarding class, the default LSP is used. If the packet is a VPLS multicast/broadcast packet and the user did not explicitly specify the LSP to use under the **config>service>sdp>class-forwarding>multicast-lsp** context, then the default LSP is used.

VLL service packets are forwarded based on their forwarding class only if shared queuing is enabled on the ingress SAP. Otherwise, the class-forwarding command has no effect and all packets of the VLL will be forwarded on one LSP of the SDP.

The **no** form of the command deletes the configuration and the SDP reverts back to forwarding service packets based on the hash algorithm used for LAG and ECMP.

Default **no class-forwarding** — Packets of a service bound to this SDP will be forwarded based on the hash algorithm used for LAG and ECMP.

Parameters **default-lsp** *lsp-name* — Specifies the default LSP for the SDP. This LSP name must already exist and must have been associated with this SDP using the **config>service>sdp>lsp** command. The default LSP is used to forward packets when there is no available LSP matching the packet's forwarding class. This is possible if the LSP associated with the packet's forwarding class is down, or if the LSP does not exist. The default LSP is also used to forward VPLS service multicast/broadcast packets in the absence of a user configuration indicating an explicit association to one of the SDP LSPs.

Note that when the default LSP is down, the SDP is also brought down.

fc

Syntax **fc** {**be** | **l2** | **af** | **l1** | **h2** | **ef** | **h1** | **nc**} **lsp** *lsp-name*
no fc {**be** | **l2** | **af** | **l1** | **h2** | **ef** | **h1** | **nc**}

Context config>service>sdp>forwarding-class

Description This command makes an explicit association between a forwarding class and an LSP. The LSP name must exist and must have been associated with this SDP using the command config>service>sdp>lsp. Multiple forwarding classes can be associated with the same LSP. However, a forwarding class can

only be associated with a single LSP in a given SDP. All subclasses will be assigned to the same LSP as the parent forwarding class.

Default **none**

Parameters **lsp** *lsp-name* — Specifies the RSVP or static LSP to use to forward service packets which are classified into the specified forwarding class.

multicast-lsp

Syntax **multicast-lsp** *lsp-name*
no multicast-lsp

Context config>service>sdp>forwarding-class

Description This command specifies the RSVP or static LSP in this SDP to use to forward VPLS multicast and broadcast packets. The LSP name must exist and must have been associated with this SDP using the command config>service>sdp>lsp. In the absence of an explicit configuration by the user, the default LSP is used.

Default **default-lsp-name**

far-end

Syntax **far-end** *ip-address*
no far-end

Context config>service>sdp

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

The far-end IP address must be explicitly configured. The destination IP address must be a SR-Series system IP address.

If the SDP uses GRE for the destination encapsulation, the *ip-address* is checked against other GRE SDPs to verify uniqueness. If the *ip-address* is not unique within the configured GRE SDPs, an error is generated and the *ip-address* is not associated with the SDP. The local SR-Series may not know whether the *ip-address* is actually a system IP interface address on the far end SR-Series.

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

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

The **no** form of this command removes the currently configured destination IP address for the SDP. The *ip-address* parameter is not specified and will generate an error if used in the **no far-end** command. The SDP must be administratively disabled using the **config service sdp shutdown**

command before the **no far-end** command can be executed. Removing the far end IP address will cause all *lsp-name* associations with the SDP to be removed.

Default **none**

Parameters *ip-address* — The system address of the far-end 7750 SR for the SDP in dotted decimal notation.

ldp

Syntax [**no**] **ldp**

Context config>service>sdp

Description This command enables LDP-signaled LSP's on MPLS-encapsulated SDPs.

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

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

Default **no ldp** (disabled)

lsp

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

Context config>service>sdp

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

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

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

Each LSP has a preference value defined that is derived from its active path or the LSP itself. When multiple LSPs are defined on an SDP, the LSP with the highest preference will be used for packets transmitted to the far-end 7750 SR.

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

shutdown, causing the association with the SDP to be operationally down (the LSP will not be used by the SDP).

LSP SDPs also require that the TLDP signaling is specified and the SDP Keepalive parameter must be enabled and and not timed out.

Up to 16 LSP names can be entered on a single command line.

See **MPLS Configuration Commands** and **RSVP Configuration Commands** for CLI syntax and command usage.

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

Default No LSPs names are defined.

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

metric

Syntax **metric** *metric*
no metric

Context config>service>sdp

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

Parameters *metric* — specifies the SDP metric.

Values 0 — 65535

signaling

Syntax **signaling** {**off** | **tldp**}

Context config>service>sdp

Description This command specifies the signaling protocol used to obtain the ingress and egress labels in frames transmitted and received on the SDP. When signaling is *off* then labels are manually configured when the SDP is bound to a service. The signalling value can only be changed while the administrative status of the SDP is down.

The **no** form of this command is not applicable. To modify the signaling configuration, the SDP must be administratively shut down and then the signaling parameter can be modified and re-enabled.

Default **tldp**

Parameters *off* — Ingress and egress signal auto-labeling is not enabled. If this parameter is selected, then each service using the specified SDP must manually configure VPN labels. This configuration is independent of the SDP's transport type, GRE, MPLS (RSVP or LDP).

tldp — Ingress and egress signalling autolabelling is enabled.

vlan-vc-etype

Syntax **vlan-vc-etype** *0x0600..0xffff*
no vlan-vc-etype [*0x0600..0xffff*]

Context config>service>sdp

Description This command configures the VLAN VC EtherType.

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

Default **no vlan-vc-etype**

Parameters *0x0600..0xffff* — Specifies a valid VLAN etype identifier.

path-mtu

Syntax **path-mtu** *bytes*
no path-mtu

Context config>service>sdp

Description This command configures the Maximum Transmission Unit (MTU) in bytes that the Service Distribution Point (SDP) can transmit to the far-end 7750 SR router without packet dropping or IP fragmentation overriding the SDP-type default path-mtu.

The default SDP-type **path-mtu** can be overridden on a per SDP basis.

Dynamic maintenance protocols on the SDP like RSVP may override this setting.

If the physical **mtu** on an egress interface or PoS channel indicates the next hop on an SDP path cannot support the current **path-mtu**, the operational **path-mtu** on that SDP will be modified to a value that can be transmitted without fragmentation.

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

Default The default **path-mtu** defined on the system for the type of SDP is used.

SDP Keepalive Commands

keep-alive

Syntax **keepalive**

Context config>service>sdp

Description Context for configuring SDP connectivity monitoring keepalive messages for the SDP ID.

SDP-ID keepalive messages use SDP Echo Request and Reply messages to monitor SDP connectivity. The operating state of the SDP is affected by the keepalive state on the SDP-ID. SDP Echo Request messages are only sent when the SDP-ID is completely configured and administratively up. If the SDP-ID is administratively down, keepalives for that SDP-ID are disabled. SDP Echo Requests (when sent for keepalive messages) are always sent with the *originator-sdp-id*. All SDP-ID keepalive SDP Echo Replies are sent using generic IP/GRE OAM encapsulation.

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

A set of event counters track the number of keepalive requests sent, the size of the message sent, non-error replies received and error replies received. A keepalive state value is kept indicating the last response event. A keepalive state timestamp value is kept indicating the time of the last event. With each keepalive event change, a log message is generated indicating the event type and the timestamp value.

The table below describes keepalive interpretation of SDP Echo Reply response conditions and the effect on the SDP ID operational status.

Result of Request	Stored Response State	Operational State
keepalive request timeout without reply	Request Timeout	Down
keepalive request not sent due to non-existent *orig-sdp-id*[a]	Orig-SDP Non-Existent	Down
keepalive request not sent due to administratively down *orig-sdp-id*	Orig-SDP Admin-Down	Down
keepalive reply received, invalid origination-id	Far End: Originator-ID Invalid	Down
keepalive reply received, invalid responder-id	Far End: Responder-ID Error	Down
keepalive reply received, No Error	Success	Up (If no other condition prevents)

a. This condition should not occur.

hello-time

Syntax **hello-time** *seconds*
no hello-time

Context config>service>sdp>keep-alive

Description Configures the time period between SDP keepalive messages on the SDP-ID for the SDP connectivity monitoring messages.

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

Default **hello-time 10** — 10 seconds between keepalive messages

seconds — The time period in seconds between SDP keepalive messages, expressed as a decimal integer.

Values 1 — 3600

hold-down-time

Syntax **hold-down-time** *seconds*
no hold-down-time

Context config>service>sdp>keep-alive

Description Configures the minimum time period the SDP will remain in the operationally down state in response to SDP keepalive monitoring.

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

When an SDP keepalive response is received that indicates an error condition or the **max-drop-count** keepalive messages receive no reply, the *sdp-id* will immediately be brought operationally down. If a keepalive response is received that indicates the error has cleared, the *sdp-id* will be eligible to be put into the operationally up state only after the **hold-down-time** interval has expired.

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

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

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

Values 0 — 3600

max-drop-count

Syntax **max-drop-count** *count*
no max-drop-count

Context config>service>sdp>keep-alive

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

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

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

Default **max-drop-count 3** — Up to 3 SDP keepalive messages requests can fail to be sent or replies missed before the SDP is brought down.

Parameters *count* — The number of consecutive SDP keepalive requests that are failed to be sent or replies missed, expressed as a decimal integer.

Values 1 — 5

message-length

Syntax **message-length** *octets*
no message-length

Context config>service>sdp>keepalive

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

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

Default 0 — The message length should be equal to the SDP's operating path MTU as configured in the **path-mtu** command.

If the default size is overridden, the actual size used will be the smaller of the operational SDP-ID Path MTU and the size specified.

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

Values 40 — 9198

timeout

Syntax **timeout** *timeout*
no timeout

Context config>service>sdp>keep-alive

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

Default **5**

Parameters *timeout* — The timeout time, in seconds.

Values 1 — 10

Show Commands

Show Service Commands

customer

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

Context show>service

Description 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.
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-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 : VPLS Customer
Phone       : (567) 555-1212

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

Customer-ID : 274
Contact     : Mssrs. Beaucoup
Description : ABC Company
Phone       : 650 123-4567

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

------------------------------------------------------
Total Customers : 8
------------------------------------------------------------
*A: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 1
I. Sched Pol: SLA1
E. Sched Pol: (Not Specified)
-------------------------------------------------------------------------------
Service Association
-------------------------------------------------------------------------------
No Service Association Found.
===============================================================================
*A:ALA-12#
```

sdp

Syntax **sdp** [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]

Context show>service

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 desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Opr MTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
IP address	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Adm Admin State	Specifies the desired state of the SDP.
Opr Oper State	Specifies the operating state of the SDP.
Deliver	Specifies the type of delivery used by the SDP: GRE or MPLS.
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.
Hello Msg Len	Specifies the length of the SDP Echo Request messages transmitted on this SDP.
Hello Timeout	Specifies the number of seconds to wait for an SDP echo response message before declaring a timeout.
Unmatched Replies	Specifies the number of SDP unmatched message replies timer expired.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
TX Hello Msgs	Specifies the number of SDP echo request messages transmitted since the keepalive was administratively enabled or the counter was cleared.

Label	Description
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
===============================================================================
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:ALA-12# show service sdp 8
===============================================================================
Service Destination Point (Sdp Id : 8)
===============================================================================
SdpId    Adm MTU   Opr MTU   IP address       Adm  Opr         Deliver Signal
-------------------------------------------------------------------------------
8        4462      4462      10.10.10.104     Up   Dn NotReady MPLS    TLDP
===============================================================================
*A:ALA-12#
===============================================================================
Service Destination Point (Sdp Id : 8) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 8  -(10.10.10.104)
-------------------------------------------------------------------------------
Description           : MPLS-10.10.10.104
SDP Id               : 8
Admin Path MTU       : 0                    Oper Path MTU      : 0
Far End              : 10.10.10.104         Delivery           : MPLS
Admin State          : Up                   Oper State         : Down
Flags                : SignalingSessDown TransportTunnDown
Signaling            : TLDP                 VLAN VC Etype      : 0x8100
Last Status Change   : 02/01/2007 09:11:39  Adv. MTU Over.     : No
Last Mgmt Change     : 02/01/2007 09:11:46

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

Associated LSP LIST :
Lsp Name             : to-104
Admin State          : Up                   Oper State         : Down
Time Since Last Tran*: 01d07h36m
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:ALA-12#
```

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.
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#
```

service-using

Syntax **service-using** [**epipe**] [**ies**] [**vpls**] [**vprn**] [**mirror**] [**apipe**] [**fpipe**] [**ipipe**] **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 **epipe —** Displays matching Epipe services.

ies — Displays matching IES instances.

vpls — Displays matching VPLS instances.

vprn — Displays matching VPRN services.

mirror — Displays matching mirror services.

apipe — Displays matching Apipe services.

fpipe — Displays matching Fpipe services.

ipipe — Displays matching Ipipe 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 a 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
-------------------------------------------------------------------------------
Matching Services : 3
-------------------------------------------------------------------------------
===============================================================================
*A:ALA-12#
```

```
*A:ALA-12# show service service-using epipe
===============================================================================
Services [epipe]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
6            Epipe     Up     Up         6                 09/22/2006 23:05:58
7            Epipe     Up     Up         6                 09/22/2006 23:05:58
8            Epipe     Up     Up         3                 09/22/2006 23:05:58
103          Epipe     Up     Up         6                 09/22/2006 23:05:58
-------------------------------------------------------------------------------
Matching Services : 4
-------------------------------------------------------------------------------
===============================================================================
*A:ALA-12#

de14# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
1            uVPLS     Up     Up         1                 10/26/2006 15:44:57
2            Epipe     Up     Down       1                 10/26/2006 15:44:57
10           mVPLS     Down   Down       1                 10/26/2006 15:44:57
11           mVPLS     Down   Down       1                 10/26/2006 15:44:57
100          mVPLS     Up     Up         1                 10/26/2006 15:44:57
101          mVPLS     Up     Up         1                 10/26/2006 15:44:57
102          mVPLS     Up     Up         1                 10/26/2006 15:44:57
999          uVPLS     Down   Down       1                 10/26/2006 16:14:33
-------------------------------------------------------------------------------
Matching Services : 8
-------------------------------------------------------------------------------
de14#
```

VLL Services

In This Chapter

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

Topics in this section include:

Ethernet Pipe (Epipe) Services

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

Topics in this section include:

Epipe Service Overview

An Epipe service is Alcatel-Lucent's implementations of an Ethernet VLL based on the IETF "Martini Drafts" (draft-martini-l2circuit-trans-mpls-08.txt and draft-martini-l2circuit-encapmpls-04.txt) and the IETF Ethernet Pseudo-wire Draft (draft-so-pwe3-ethernet-00.txt).

An Epipe service is a layer 2 point-to-point service where the customer data is encapsulated and transported across a service provider's IP or MPLS network. An Epipe service is completely transparent to the subscriber's data and protocols. The 7750 SR Epipe service does not perform any MAC learning. A local Epipe service consists of two SAPs on the same node, whereas a distributed Epipe service consists of two SAPs on different nodes. SDPs are not used in local Epipe services.

Each SAP configuration includes a specific port/channel on which service traffic enters the 7750 SR from the customer side (also called the access side). Each port is configured with an encapsulation type. If a port is configured with an IEEE 802.1Q (referred to as dot1q) encapsulation, then a unique encapsulation value (ID) must be specified.

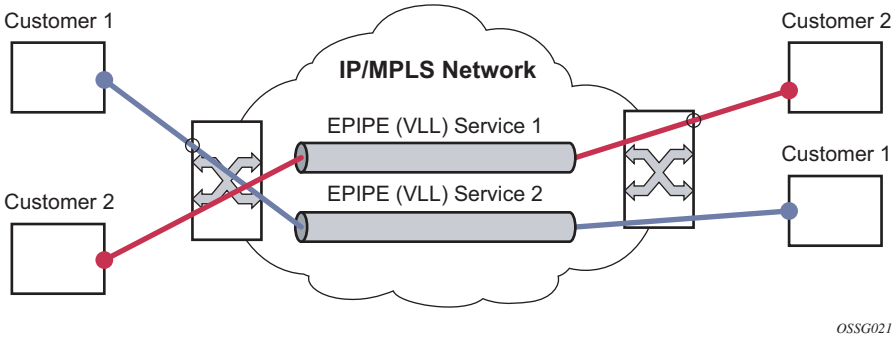

Figure 1: Epipe/VLL Service

Ethernet Interworking VLL

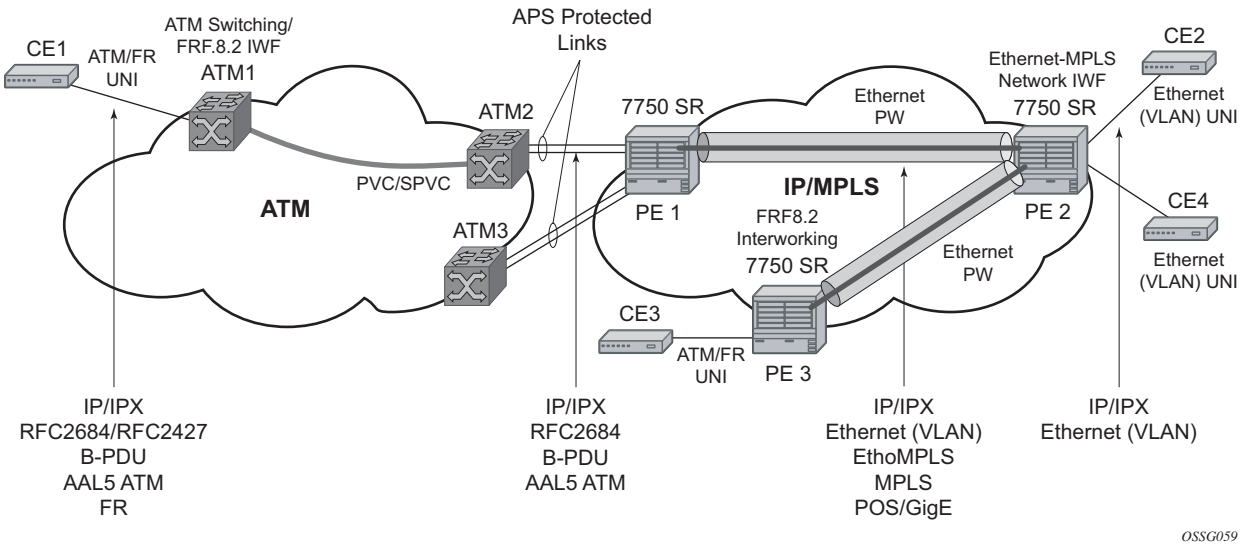

Figure 2: Application of Ethernet Interworking VLL Example

Figure 15 provides an example of an Ethernet interworking VLL. The Ethernet interworking VLL provides a point-to-point Ethernet VLL service between Frame-Relay-attached users, ATM attached users, and Ethernet-attached users across an IP/MPLS packet switched network. It effectively provides ATM and FR bridged encapsulation termination on the existing Epipe service of the 7750 SR.

The following connectivity scenarios are supported:

- A Frame Relay or ATM user connected to a ATM network communicating with a Ethernet user connected to a 7750 PE node on a IP/MPLS network.
- A Frame Relay or ATM user connected to 7750 SR PE node communicating with an Ethernet user connected to a 7750 PE node on a IP/MPLS network. This feature supports local cross-connecting when these users are attached to the same 7750 PE node.

Users attach over an ATM UNI with RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, tagged/untagged bridged Ethernet PDUs, a FR UNI using RFC 2427, *Multiprotocol Interconnect over Frame Relay*, tagged/untagged bridged Ethernet PDUs, or an Ethernet tagged/untagged UNI interface. However, the VCI/VPI and the data-link control layer (DLCI) are the identifiers of the SAP in the case of ATM and FR respectively and the received tags are transparent to the service and are thus preserved.

The Ethernet pseudowire is established using Targeted LDP (TLDP) signaling and can use the **ether** or **vlan** VC types on the SDP. The SDP can be either an MPLS or GRE type.

ATM VLL (Apipe) Services

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

Topics in this section include:

ATM VLL For End-to-End ATM Service

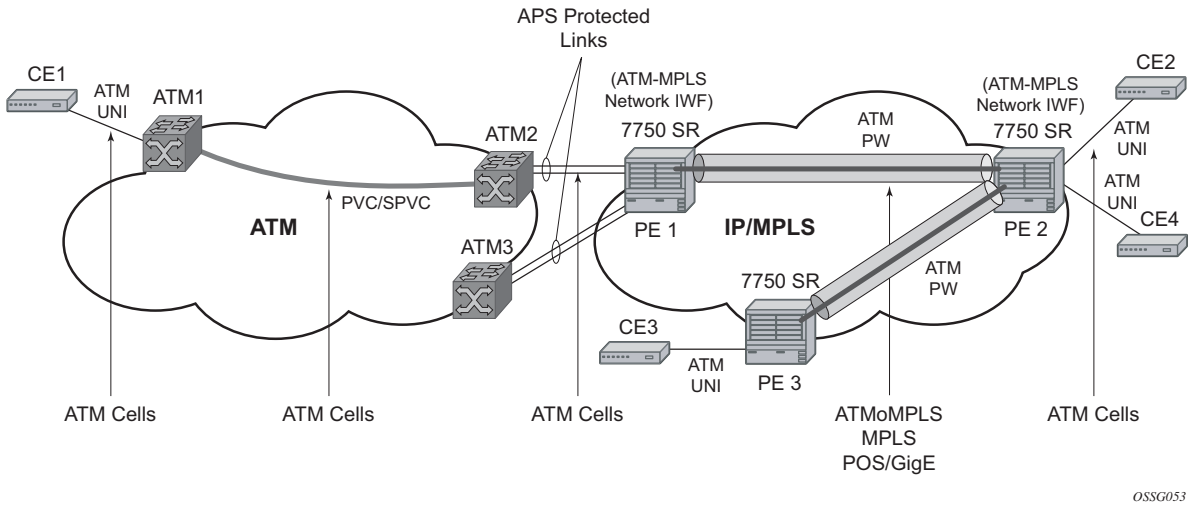

Figure 3: ATM VLL for End-to-End ATM Service

ATM VLLs (Apipe) provide a point-to-point ATM service between users connected to 7750 SR nodes on an IP/MPLS network. Users are either directly connected to a 7750 PE or through an ATM access network. In both cases, an ATM PVC (for example, a virtual channel (VC) or a virtual path (VP)) is configured on the 7750 PE. This feature supports local cross-connecting when users are attached to the same 7750 PE node. VPI/VCI translation is supported in the ATM VLL.

PE1, PE2, and PE3 receive standard UNI/NNI cells on the ATM Service Access Point (SAP) that are then encapsulated into a pseudowire packet using the N:1 cell mode encapsulation or AAL5 SDU mode encapsulation according to draft-ietf-pwe3-atm-encap-xx.txt, *Encapsulation Methods for Transport of ATM Over MPLS Networks*. When using N:1 cell mode encapsulation, cell concatenation into a pseudowire packet is supported. In this application, the setup of both VC and VP level connections are supported.

The ATM pseudowire is initiated using Targeted LDP (TLDP) signaling as specified in draft-ietf-pwe3-control-protocol-xx.txt, *Pseudowire Setup and Maintenance using LDP.* The SDP can be an MPLS or a GRE type.

ATM Virtual Trunk Over IP/MPLS Packet-Switched Network

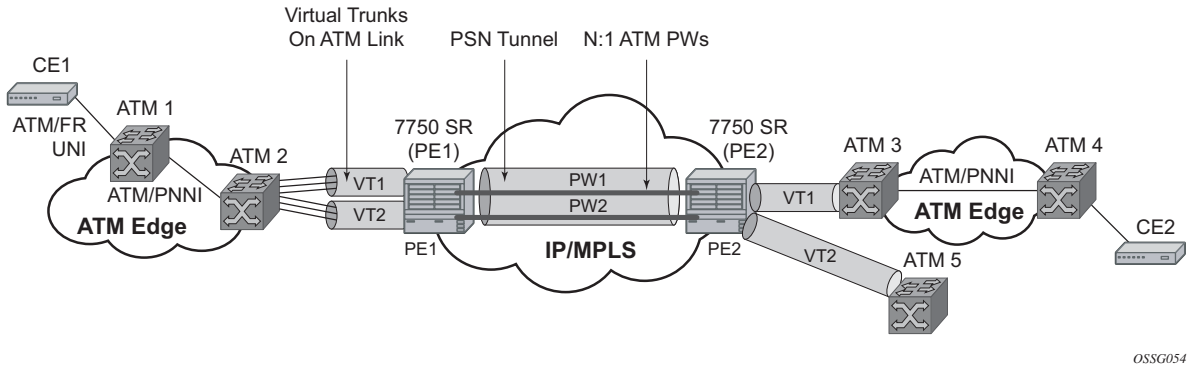

Figure 4: VT Application Example

ATM virtual trunk (VT) implements a transparent trunking of user and control traffic between two ATM switches over an ATM pseudowire. Figure 17 depicts ATM 2 and ATM 3 switches that appear as if they are directly connected over an ATM link. Control traffic includes PNNI signaling and routing traffic.

The virtual trunk (VT) SAP on a 7750 PE is identified by a tuple (port, VPI-range) meaning that all cells arriving on the specified port within the specified VPI range are fed into a single ATM pseudowire for transport across the IP/MPLS network. Note that a user can configure the whole ATM port as a VT and does not need to specify a VPI range. No VPI/VCI translation is performed on ingress or egress. Cell order is maintained within a VT. Note that, as a special case, the two ATM ports could be on the same PE node.

By carrying all cells from all VPIs making up the VT in one pseudowire, a solution is provided that is both robust, for example no black holes on some VPIs but not others, as well as operationally efficient since the entire VT can be managed as a single entity from the Network Manager (single point for configuration, status, alarms, statistics, etc.).

ATM virtual trunks use PWE3 N:1 ATM cell mode encapsulation to provide a cell-mode transport, supporting all AAL types, over the MPLS network. Cell concatenation on a pseudowire packet is supported. The SDP can be of an MPLS or a GRE type.

The ATM PW is initiated using Targeted LDP (TLDP) signaling (defined in draft-ietf-pwe3-control-protocol-xx.txt, *Pseudowire Setup and Maintenance using LDP*). In this application, there is no ATM signaling on the 7750 gateway nodes since both endpoints of the MPLS network are configured by the network operator. ATM signaling between the ATM nodes is passed transparently over the VT (along with user traffic) from one ATM port on a 7750 PE to another ATM port on a remote (or the same) 7750 SR PE.

Traffic Management Support

Ingress Network Classification

Classification is based on the EXP value of the pseudowire label and EXP-to-FC mapping is determined by the network ingress QoS policy.

Ingress Queuing and Shaping on the IOM

Each SAP of an ATM VLL has an associated single ingress service queue on the IOM. The default QoS policy configures this queue to have CIR=0 and PIR=line rate. Other QoS policies can be applied if they specify a single service queue. Applying a non-default QoS policy allows the CIR/PIR of the incoming traffic to be controlled, regardless of whether ATM policing is configured, and provides queuing and shaping to smooth traffic flows on the ingress of the network.

Egress Queuing and Shaping on the IOM

Each SAP of an ATM VLL has a single associated egress service queue on the IOM. The default QoS policy configures this queue to have CIR=0 and PIR=line rate. Other QoS policies can be applied if they specify a single service queue. Applying a non-default QoS policy allows the CIR/PIR of the outgoing traffic to be controlled, regardless of whether ATM shaping is configured.

Egress Shaping/Scheduling

Each SAP of an ATM VLL has an egress ATM traffic descriptor associated with it. The default traffic descriptor has service category UBR with zero MIR, resulting in endpoints associated with this descriptor being scheduled at the lowest priority on the ATM MDA. Egress traffic may be shaped or scheduled, depending on the configuration of the egress ATM traffic descriptor associated with the SAP. Table 3 provides details of how the different service categories and shaping settings affect egress transmission rates.

Shaping applies to CBR, rtVBR and nrtVBR service categories and results in cells being transmitted in such a way as to satisfy a downstream ATM UPC function. In particular, the transmission rate will be limited (in the case of CBR, there is a hard limit of PIR, while rtVBR/nrtVBR will transmit at SIR with short (constrained by MBS) bursts of up to PIR), and the inter-cell gap will also be controlled.

Service category UBR and rtVBR are scheduled on the WRR scheduler with the configured rates (MIR for UBR+) determining the weight applied to the flow. Weights are between 1 and 255 and are determined by a formula applied to the configured rate. UBR flows (i.e., those with no MIR) receive a weight of 1 and the maximum weight of 255 is reached by flows with configured rates of

around 8 Mbps. Scheduling does not apply a limit to the transmission rate; the available port bandwidth is shared out by the scheduler according to the weight, so if other flows are quiescent, a given flow may burst up to port bandwidth.

Shaping and scheduling of egress ATM VLL traffic is performed entirely at the ATM layer and is therefore not forwarding-class-aware. If the offered rate is greater than can be transmitted towards the customer (either because the shaping rate limits transmission or because the SAP does not receive sufficient servicing in the weighed round-robin used for scheduled SAPs), the per-VC queue will back up and this will trigger the congestion control mechanisms in the MDA queues or in the IOM service egress queues associated with the SAP. For AAL5 SDU VLLs, these discards occur at the AAL5 SDU level. For N-to-1 VLLs, these discards occur at the level of the cell or a block of cells when cell concatenation is enabled.

Table 1: Behavior and Relative Priorities

Flow type	Transmission rate	Priority
shaped CBR	Limited to configured PIR	Strict priority over all other traffic
shaped rtVBR	Limited to configured SIR, but with bursts up to PIR within MBS	Strict priority over all but shaped CBR
shaped nrtVBR	Limited to configured SIR, but with bursts up to PIR within MBS	Strict priority over all scheduled traffic
scheduled nrtVBR	Weighted share (according to SIR) of port bandwidth remaining after shaped traffic has been exhausted	In the same WRR scheduler as UBR+ and UBR
scheduled UBR+	Weighted share (according to MIR) of port bandwidth remaining after shaped traffic has been exhausted	In the same WRR scheduler as nrtVBR and UBR
scheduled UBR	Weighted share (with weight of 1) of port bandwidth remaining after shaped traffic has been exhausted	In the same WRR scheduler as nrtVBR and UBR+

Frame Relay VLL (Fpipe) Services

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

Topics in this section include:

Frame Relay VLL

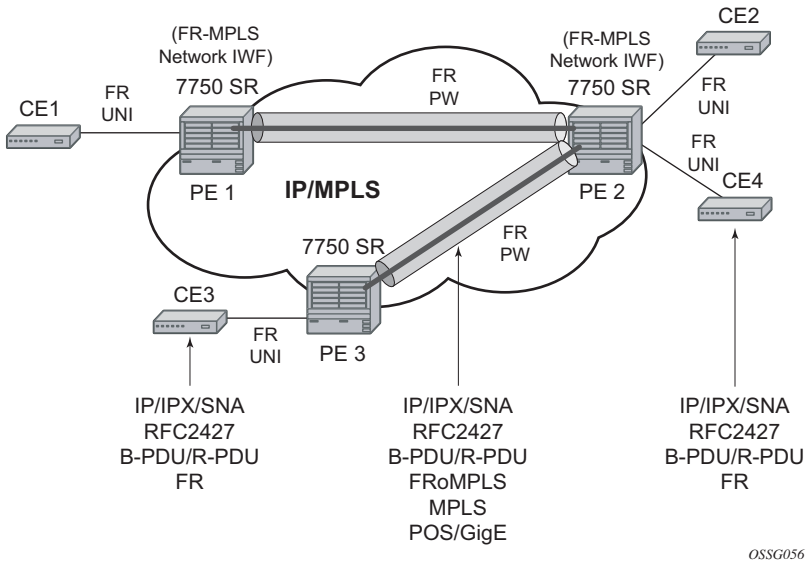

Figure 5: Application of a Frame Relay VLL Example

Figure 18 depicts an application of a Frame Relay VLL. The Frame Relay VLL (Fpipe) provides a point-to-point Frame Relay service between users connected to 7750 nodes on the IP/MPLS network. Users are connected to the 7750 PE nodes using Frame Relay PVCs. PE1, PE2, and PE3

receive a standard Q.922 Core frame on the Frame Relay SAP and encapsulate it into a pseudowire packet according to the 1-to-1 Frame Relay encapsulation mode in draft-ietf-pwe3-frame-relay-xx.txt, *Encapsulation Methods for Transport of Frame Relay Over MPLS Networks*. The 7750 Frame Relay VLL feature supports local cross-connecting when the users are attached to the same 7750 PE node.

The FR PW is initiated using Targeted LDP (TLDP) signaling as specified in draft-ietf-pwe3-control-protocol-xx.txt, *Pseudowire Setup and Maintenance using LDP*. The SDP can be an MPLS or a GRE type.

Frame Relay-to-ATM Interworking (FRF.5) VLL

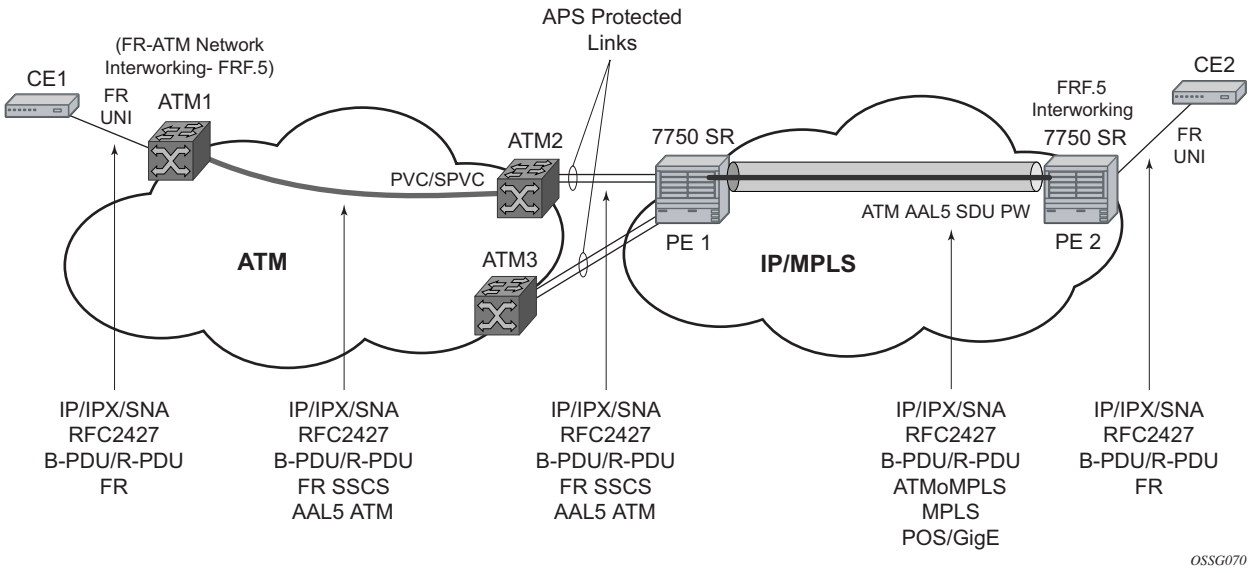

Figure 6: Frame Relay-to-ATM Network Interworking (FRF.5) VLL

Figure 19 provides an example of a point-to-point Frame Relay service between users where one user is connected to an existing ATM network, the other to a 7750 SR PE node on an IP/MPLS network.

This VLL is realized using an ATM AAL5 SDU PW between the 7750 SR PE nodes. It is configured by adding a FR SAP to an apipe service using vc-type atm-sdu. The 7750 SR PE2 node performs an FRF.5 interworking function to interwork the ingress and egress data paths in addition to the operations required in an FR and an ATM VLL.

The PW is initiated using Targeted LDP signaling as specified in draft-ietf-pwe3-control-protocol-xx.txt. The SDP can be of an MPLS or a GRE type.

Traffic Management Support

Frame Relay Traffic Management

Traffic management of Frame Relay VLLs is achieved through the application of ingress and egress QoS policies to SAPs like other Frame Relay SAPs. No queuing occurs on the MDA; all queuing, policing and shaping occurs on the IOM and, as a result, traffic management is forwarding-class-aware. Forwarding classes may be determined by inspecting the DSCP marking of contained IP packets (for example) and this will determine both the queuing and the EXP bit setting of packets on a Frame Relay VLL.

Ingress SAP Classification and Marking

DE=0 frames are subject to the CIR marking algorithm in the IOM queue. Drop preference for these packets will follow the state of the CIR bucket associated with the ingress queue. The value is marked in the drop preference bit of the internal header and into the DE bit in the Q.922 frame header. DE=1 frames are classified into “out-of-profile” state and are not be overwritten by the CIR marking in the ingress IOM queue. The drop preference is set to high.

Egress Network EXP Marking

FC-to-EXP mapping is as per the Network Egress QoS policy. Marking of the EXP field in both label stacks is performed.

Ingress Network Classification

Classification is based on the EXP value of the pseudowire label and EXP-to-FC mapping is as per Network Ingress QoS policy.

IP Interworking VLL (Ipipe) Services

Ipipe VLL

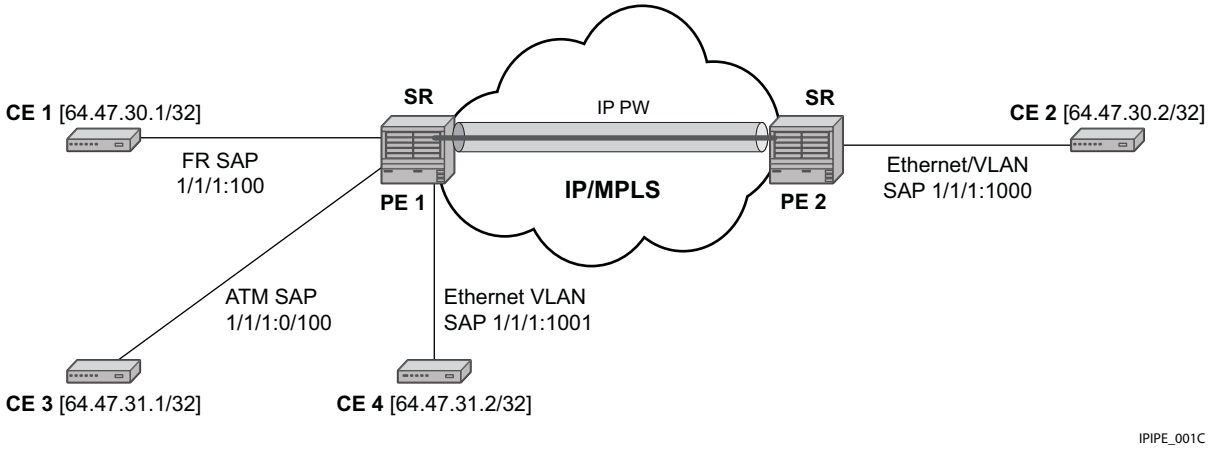

Figure 7: IP Interworking VLL (Ipipe)

Figure 20 provides an example of IP connectivity between a host attached to a point-to-point access circuit (FR, ATM, PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts appear to be on the same LAN segment. This feature enables service interworking between different link layer technologies. A typical use of this application is in a Layer 2 VPN when upgrading a hub site to Ethernet while keeping the spoke sites with their existing Frame Relay or ATM IPv4 routed encapsulation.

The ATM SAP carries the IPv4 packet using RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5, VC-Mux or LLC/SNAP routed PDU encapsulation.* The Frame Relay SAP makes use of RFC 2427, *Multiprotocol Interconnect over Frame Relay,* routed PDU encapsulation

of an IPv4 packet. A PPP interface makes use of RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*, PPP IPCP encapsulation of an IPv4 packet. A Cisco HDLC SAP uses the routed IPv4 encapsulation. The pseudowire uses the IP Layer 2 transport pseudowire encapsulation type.

Note that the Ipipe is a point-to-point Layer 2 service. All packets received on one SAP of the Ipipe will be forwarded to the other SAP. No IP routing of customer packets occurs.

IP Interworking VLL Datapath

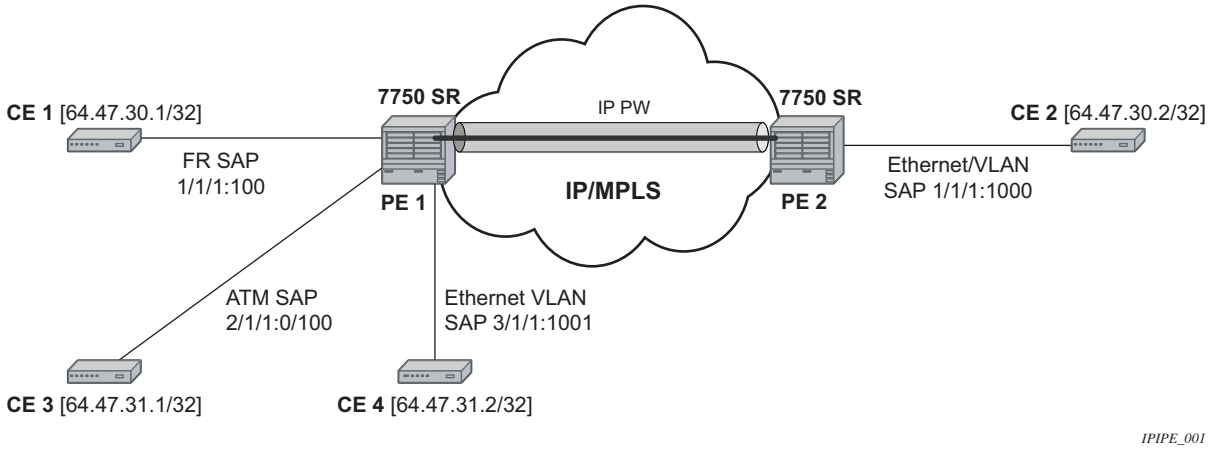

Figure 8: IP Interworking VLL Datapath

In order to be able to forward IP packets between CE 1 and CE 2 in Figure 21, the PE 2 is manually configured with both CE 1 and CE 2 IP addresses. These are host addresses and are entered in the /32 format. PE 2 maintains an ARP cache context for each IP interworking VLL. PE 2 responds to ARP request messages received on the Ethernet SAP. PE 2 responds with the Ethernet SAP configured MAC address as a proxy for any ARP request for CE 1 IP address. PE 2 should silently discard any ARP request message received on the Ethernet SAP for an address other than that of CE 1. Likewise, PE 2 should silently discard any ARP request message with the source IP address other than that of CE 2. In all cases, PE 2 keeps track of the association of IP to MAC addresses for ARP requests it receives over the Ethernet SAP.

In order to forward unicast frames destined to CE 2, PE 2 needs to know CE 2 MAC address. When the Ipipe SAP is first configured and administratively enabled, PE2 sends an ARP request message for CE 2 MAC address over the Ethernet SAP. Until an ARP reply is received from CE2, providing CE2's MAC address, unicast IP packets destined for CE2 will be discarded at PE2. IP broadcast and IP multicast packets are sent on the Ethernet SAP using the broadcast or direct-mapped multicast MAC address.

In order to forward unicast frames destined to CE 1, PE 2 validates the MAC destination address of the received Ethernet frame. It should match that of the Ethernet SAP. It then removes the Ethernet header and encapsulates the IP packet directly into a PW without a control word. PE 1 removes the PW encapsulation and forwards the IP packet over the Frame Relay SAP using RFC 2427, *Multiprotocol Interconnect over Frame Relay,* routed PDU encapsulation.

In order to forward unicast packets destined to CE1, PE2 validates the MAC destination address of the received Ethernet frame. If the IP packet is unicast, the MAC destination must match that of the Ethernet SAP. If the IP packet is multicast or broadcast, the MAC destination address must be an appropriate multicast or broadcast MAC address. The other procedures are similar to the case of

communication between CE 1 and CE 2, except that the ATM SAP and the Ethernet SAP are cross-connected locally and IP packets do not get sent over an SDP.

A PE does not flush the ARP cache unless the SAP goes admin or operationally down. The PE with the Ethernet SAP sends unsolicited ARP requests to refresh the ARP cache every T seconds. ARP requests are staggered at an increasing rate if no reply is received to the first unsolicited ARP request. T is configurable by user through the mac-refresh CLI command.

Pseudowire Switching

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

Services with one SAP and one spoke SDP are created normally on the PE; however, the target destination of the SDP is the 7750 SR pseudowire switching node instead of what is normally the remote PE. In addition, the user configures a VLL service on the pseudowire switching node using the two SDPs.

The pseudowire switching node acts in a passive role with respect to signalling of the pseudowires. It waits until one or both of the PEs sends the label mapping message before relaying it to the other PE. This is because it needs to pass the Interface Parameters of each PE to the other.

A pseudowire switching point TLV is inserted by the switching pseudowire to record its system address when relaying the label mapping message. This TLV is useful in a few situations:

- It allows for troubleshooting of the path of the pseudowire especially if multiple pseudowire switching points exist between the two PEs.
- It helps in loop detection of the T-LDP signalling messages where a switching point would receive back a label mapping message it had already relayed.
- The switching point TLV is inserted in pseudowire status notification messages when they are sent end-to-end or from a pseudowire switching node towards a destination PE.

Pseudowire OAM is supported for the manual switching pseudowires and allows the 7750 SR pseudowire switching node to relay end-to-end pseudowire status notification messages between the two PEs. The 7750 SR pseudowire switching node can generate a pseudowire status and to send it to one or both of the PEs by including its system address in the pseudowire switching point TLV. This allows a 7750 SR PE to identify the origin of the pseudowire status notification message.

In the Figure 22, the user configures a regular Epipe VLL service in 7x50 PE1 and 7x50 PE2. These services consist each of a SAP and a spoke SPD. However, the target destination of the SDP is actually not the remote PE but the 7x50 pseudowire switching node. In addition, the user configures an epipe VLL service on the 7x50 pseudowire switching node using the two SDPs.

```
|7x50 PE1 (Epipe)|---sdp 2:10---|7x50 PW SW (Epipe)|---sdp 7:15---|7x50 PE2 (Epipe)
```

Figure 9: Pseudowire Service Switching Node

Configuration examples can be found in Configuring Two VLL Paths Terminating on T-PE2 on page 209.

Pseudowire Switching with Protection

Figure 23 illustrates the use of pseudowire redundancy to provide a scalable and resilient VLL service across multiple IGP areas in a provider network.

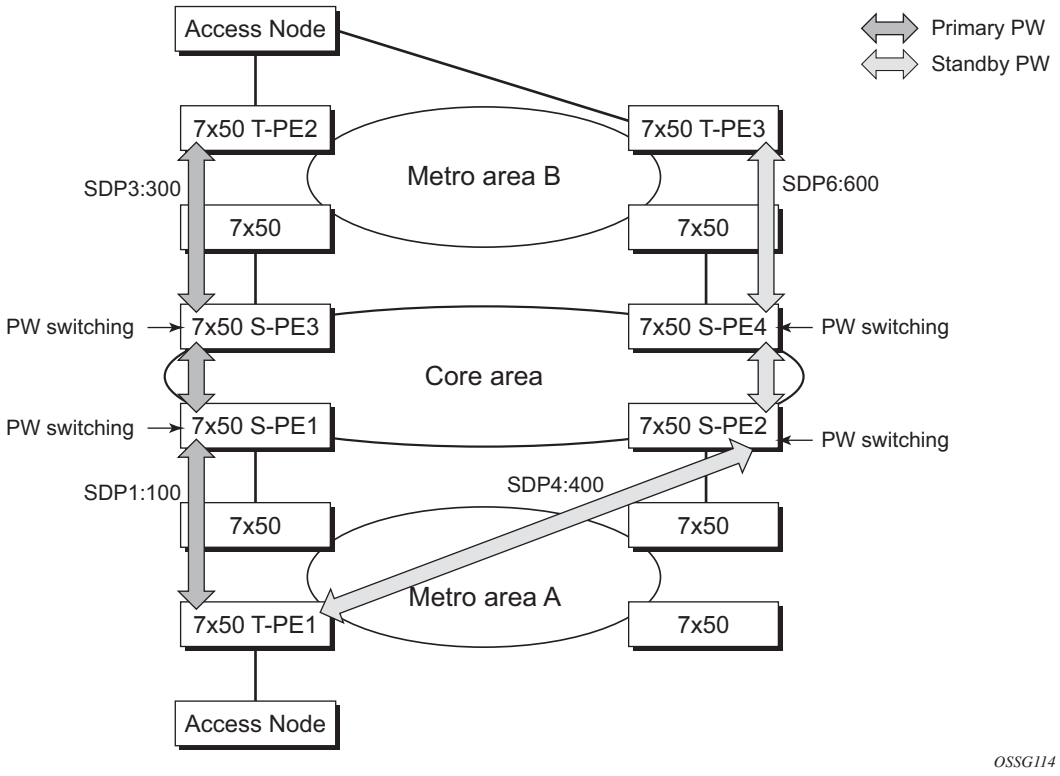

Figure 10: VLL Resilience with Pseudowire Redundancy and Switching

In the network shown in Figure 23, the user configures a VLL service, for example, an Epipe, at the ingress PE (7x50 T-PE1), at the egress PE (7x50 T-PE2), and at each pseudowire switching node (S-PE1 and S-PE2). All VLL service types, such as, Apipe, Epipe, Fpipe, and Ipipe are supported. The pseudowire switching service at S-PE is supported only between the pseudowire of the same service type. Note that the user could also configure a VPLS service at the T-PE nodes. However, the service at the switching nodes is still an Ethernet VLL service in this case.This may be added in a future release. VC-ID values are significant on a hop-by-hop basis and are unique on a per T- LDP session only.

The pseudowire switching capability can bind two spoke SDPs of different transport types, such as, RSVP LSP and LDP LSP. There is no restriction on mixing the SDP types currently supported

by VLL and VPLS services. A mix of statically and dynamically allocated PW VC labels is not supported.

Pseudowire Switching Behavior

In the network in Figure 23, PE nodes act as masters and pseudowire switching nodes act as slaves for the purpose of pseudowire signaling. This is because a switching node will need to pass the SAP interface parameters of each PE to the other.T-PE1 sends a label mapping message for the Layer 2 FEC to the peer pseudowire switching node, for example, S-PE1. It will include the SAP interface parameters, such as MTU, in the label mapping message. S-PE1 checks the FEC against the local information and if a match exists, it appends the optional pseudowire switching point TLV to the FEC TLV in which it records its system address. T-PE1 then relays the label mapping message to S-PE2. S-PE2 performs similar operation and forwards a label mapping message to T-PE2. The same procedures are followed for the label mapping message in the reverse direction, for example, from T-PE2 to T-PE1. S-PE1 and S-PE2 will effect the spoke SDP cross-connect only when both directions of the pseudowire have been signaled and matched.

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

The merging of the received T-LDP status notification message and the local status for the spoke SDPs from the service manager at a 7x50 S-PE complies with the following rules:

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

Pseudowire Switching TLV

The format of the pseudowire switching TLV is as follows:

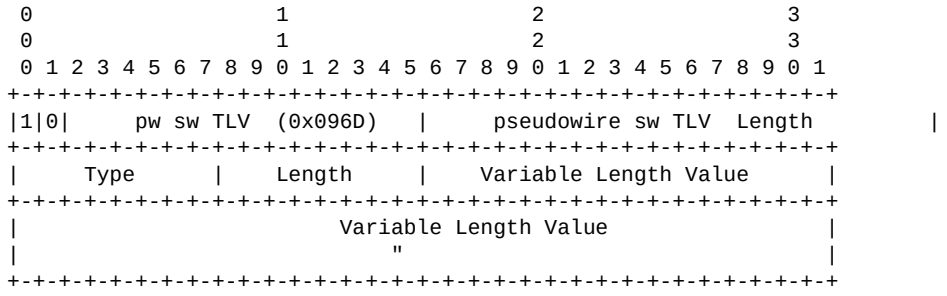

- PW sw TLV Length — Specifies the total length of all the following pseudowire switching point TLV fields in octets
- Type — Encodes how the Value field is to be interpreted.
- Length — Specifies the length of the Value field in octets.
- Value — Octet string of Length octets that encodes information to be interpreted as specified by the Type field.

Pseudowire Switching Point Sub-TLVs

Below are details specific to pseudowire switching point sub-TLVs:

- pseudowire ID of last pseudowire segment traversed — This sub-TLV type contains a pseudowire ID in the format of the pseudowire ID
- Pseudowire switching point description string — An optional description string of text up to 80 characters long.
- IP address of pseudowire switching point.
- The IP V4 or V6 address of the pseudowire switching point. This is an optional sub-TLV.
- MH VCCV capability indication.

Pseudowire Redundancy

Pseudowire redundancy provides the ability to protect a pseudowire with a pre-provisioned pseudowire and to switch traffic over to the secondary standby pseudowire in case of a SAP and/or network failure condition. Normally, pseudowires are redundant by the virtue of the SDP redundancy mechanism. For instance, if the SDP is an RSVP LSP and is protected by a secondary standby path and/or by Fast-Reroute paths, the pseudowire is also protected. However, there are a couple of applications in which SDP redundancy does not protect the end-to-end pseudowire path:

- There are two different destination 7x50 PE nodes for the same VLL service. The main use case is the provision of dual-homing of a CPE or access node to two 7x50 PE nodes located in different POPs. The other use case is the provision of a pair of active and standby BRAS nodes, or active and standby links to the same BRAS node, to provide service resiliency to broadband service subscribers.
- The pseudowire path is switched in the middle of the network and the 7x50 pseudowire switching node fails. The 7x50 supports pseudowire switching.

VLL Resilience with Two Destination PE Nodes

Figure 24 illustrates the application of pseudowire redundancy to provide Ethernet VLL service resilience for broadband service subscribers accessing the broadband service on the service provider BRAS.

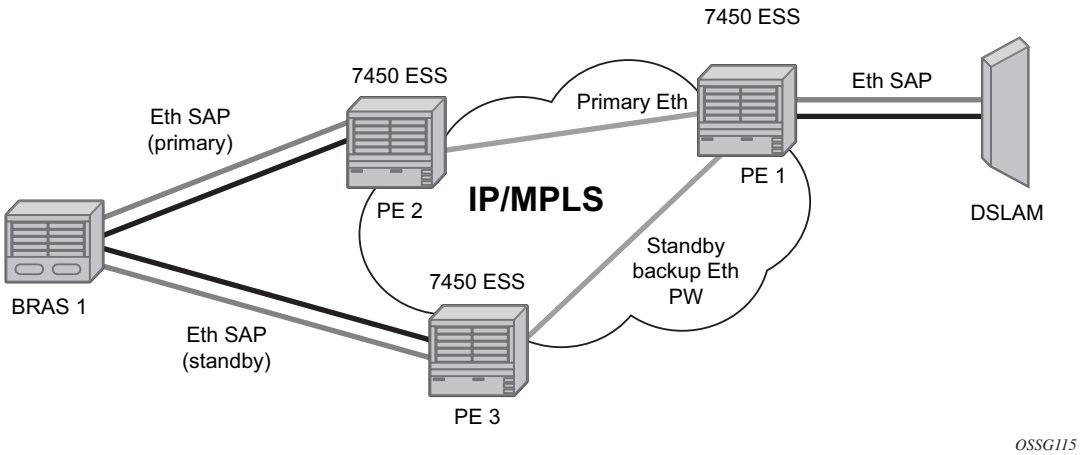

Figure 11: VLL Resilience

If the Ethernet SAP on PE2 fails, PE2 notifies PE1 of the failure by either withdrawing the primary pseudowire label it advertised or by sending a pseudowire status notification with the code set to indicate a SAP defect. PE1 will receive it and will immediately switch its local SAP to forward over the secondary standby spoke SDP. In order to avoid black holing of in-flight packets during the switching of the path, PE1 will accept packets received from PE2 on the primary pseudowire while transmitting over the backup pseudowire. Note that in the case of the FT application shown in Figure 24, this does not matter as the subscriber PPPoE session will actually timeout and will be re-established via the standby SAP on the link between PE3 and the BRAS node. However, in other applications such as those described in Access Node Resilience using MC-LAG and Pseudowire Redundancy on page 133, it will be important to minimize service outage to end users.

When the SAP at PE2 is restored, PE2 updates the new status of the SAP by sending a new label mapping message for the same pseudowire FEC or by sending pseudowire status notification message indicating that the SAP is back up. PE1 then starts a timer and reverts back to the primary at the expiry of the timer. By default, the timer is set to 0, which means PE1 reverts immediately. A special value of the timer (infinity) will mean that PE1 should never revert back to the primary pseudowire.

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

1. T-LDP peer (remote PE) node withdrew the pseudowire label.
2. T-LDP peer signaled a FEC status indicating a pseudowire failure or a remote SAP failure.
3. T-LDP session to peer node times out.
4. SDP binding and VLL service went down as a result of network failure condition such as the SDP to peer node going operationally down.

The SDP type for the primary and secondary pseudowires need not be the same. In other words, the user can protect a RSVP-TE based spoke SDP with a LDP or GRE based one. This provides the ability to route the path of the two pseudowires over different areas of the network. All VLL service types, for example, Apipe, Epipe, Fpipe, and Ipipe are supported.

The 7x50 supports the ability to configure multiple secondary standby pseudowire paths. For example, PE1 uses the value of the user configurable precedence parameter associated with each spoke SDP to select the next available pseudowire path after the failure of the current active pseudowire (whether it is the primary or one of the secondary pseudowires). The revertive operation always switches the path of the VLL back to the primary pseudowire though. There is no revertive operation between secondary paths meaning that the path of the VLL will not be switched back to a secondary pseudowire of higher precedence when the latter comes back up again.

The 7x50 supports the ability for a user-initiated manual switchover of the VLL path to the primary or any of the secondary be supported to divert user traffic in case of a planned outage such as in node upgrade procedures.

This application can make use of all types of VLL supported on the 7x50. A SAP can be configured on SONET/SDH port which is part of an APS group. However, if a SAP is configured on a MC-LAG instance, only the Epipe service type is allowed.

Access Node Resilience using MC-LAG and Pseudowire Redundancy

Figure 25 shows the use of both Multi-Chassis Link Aggregation (MC-LAG) in the access network and pseudowire redundancy in the core network to provide a resilient end-to-end VLL service to the customers.

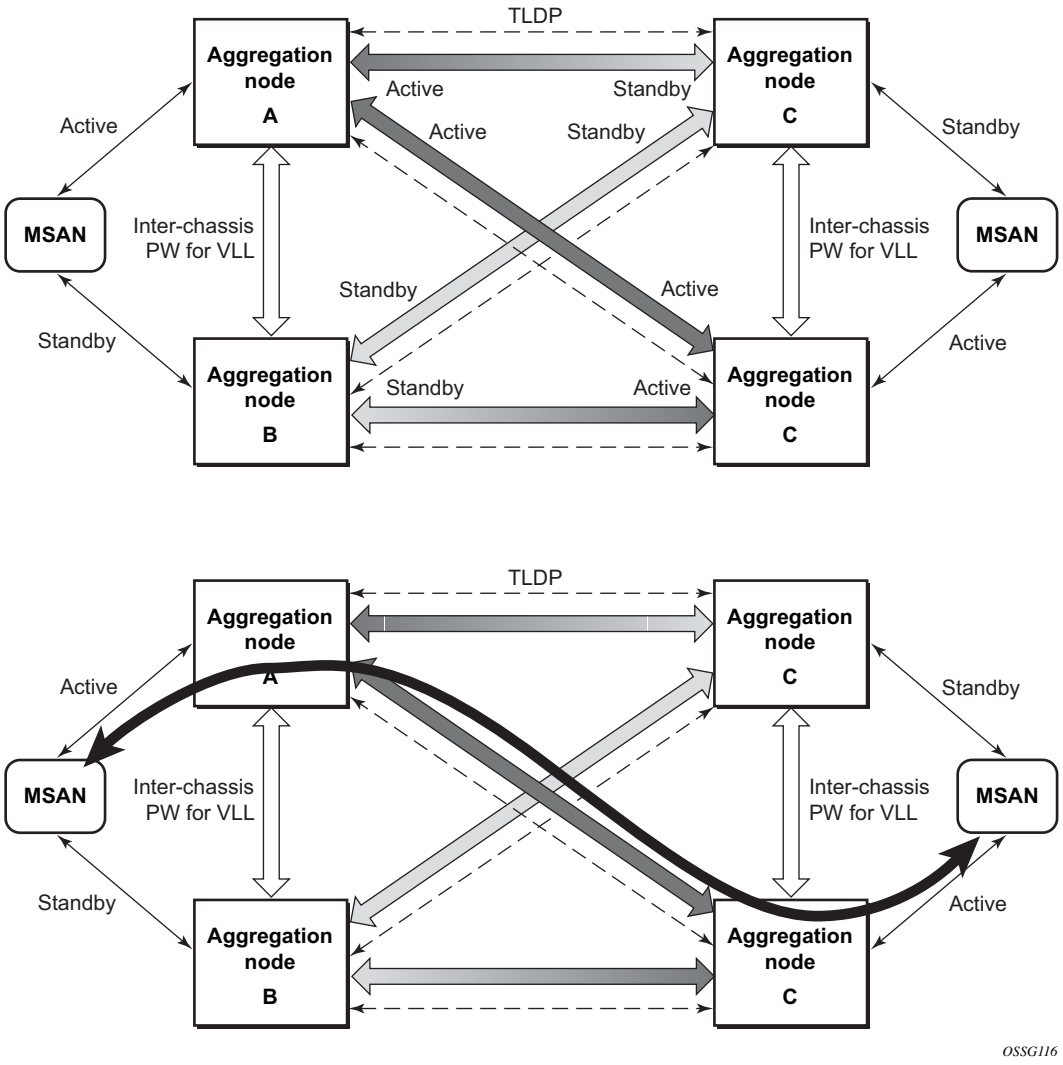

Figure 12: Access Node Resilience

In this application, a new pseudowire status bit of active or standby indicates the status of the SAP in the MC-LAG instance in the 7x50 aggregation node. All spoke SDPs are of secondary type and there is no use of a primary pseudowire type in this mode of operation. Node A is in the active state according to its local MC-LAG instance and thus advertises active status notification messages to both its peer pseudowire nodes, for example, nodes C and D. Node D performs the same operation. Node B is in the standby state according to the status of the SAP in its local MC-LAG instance and thus advertises standby status notification messages to both nodes C and D. Node C performs the same operation.

A 7x50 node selects a pseudowire as the active path for forwarding packets when both the local pseudowire status and the received remote pseudowire status indicate active status. However, a 7x50 in standby status according to the SAP in its local MC-LAG instance is capable of processing packets for a VLL service received over any of the pseudowires which are up. This is to avoid black holing of user traffic during transitions. The 7x50 standby node forwards these packets to the active node via the Inter-Chassis Backup pseudowire (ICB pseudowire) for this VLL service. An ICB is a spoke SDP used by a MC-LAG node to backup a MC-LAG SAP during transitions. The same ICB can also be used by the peer MC-LAG node to protect against network failures causing the active pseudowire to go down.

Note that at configuration time, the user specifies a precedence parameter for each of the pseudowires which are part of the redundancy set as in the application in VLL Resilience with Two Destination PE Nodes on page 130. A 7x50 node uses this to select which pseudowire to forward packet to in case both pseudowires show active/active for the local/remote status during transitions.

Only VLL service of type epipe is supported in this application. Furthermore, ICB spoke SDP can only be added to the SAP side of the VLL cross-connect if the SAP is configured on a MC-LAG instance.

More details on the behavior of this scenario can be found in Access Node Resilience using MC-LAG and Pseudowire Redundancy on page 133.

VLL Resilience for a Switched Pseudowire Path

Figure 26 illustrates the use of both pseudowire redundancy and pseudowire switching to provide a resilient VLL service across multiple IGP areas in a provider network.

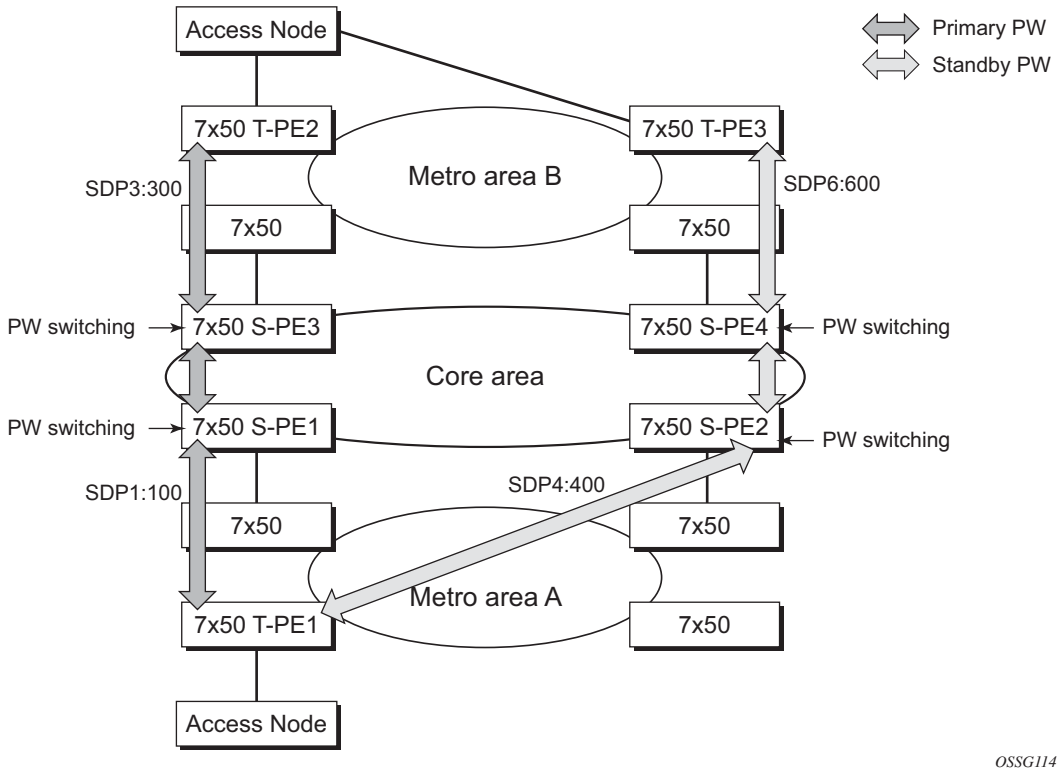

Figure 13: VLL Resilience with Pseudowire Redundancy and Switching

Pseudowire switching is a method for scaling a large network of VLL or VPLS services by removing the need for a full mesh of T-LDP sessions between the 7x50 PE nodes as the number of these nodes grows over time. In the network in Figure 4

Like in the application in VLL Resilience with Two Destination PE Nodes on page 130, 7x50 T-PE1 node switches the path of a VLL to a secondary standby pseudowire in the case of a network side failure causing the VLL binding status to be DOWN or if T-PE2 notified it that the remote SAP went down. This application requires that pseudowire status notification messages generated by either a 7x50 T-PE node or a 7x50 S-PE node be processed and relayed by the S-PE nodes.

Note that it is possible that the secondary pseudowire path terminates on the same target PE as the primary, for example, T-PE2. This provides protection against network side failures but not

against a remote SAP failure. When the target destination PE for the primary and secondary pseudowires is the same, T-PE1 will normally not switch the VLL path onto the secondary pseudowire upon receipt of a pseudowire status notification indicating the remote SAP is DOWN since the status notification is sent over both the primary and secondary pseudowires. However, the status notification on the primary pseudowire may arrive earlier than the one on the secondary pseudowire due to the differential delay between the paths. This will cause T-PE1 to switch the path of the VLL to the secondary standby pseudowire and remain there until the status notification is cleared. At that point in time, the VLL path is switched back to the primary pseudowire due to the revertive behavior operation. The path will not switch back to a secondary path when it becomes up even if it has a higher precedence than the currently active secondary path.

This application can make use of all types of VLL supported on the 7x50, for example, Apipe, Fpipe, Epipe, and Ipipe services. A SAP can be configured on SONET/SDH port which is part of an APS group. However, if a SAP is configured on a MC-LAG instance, only the Epipe service type will be allowed.

Pseudowire Redundancy at a Pseudowire Switching Node

It is possible to provide additional VLL resilience by configuring pseudowire redundancy at the S-PE nodes. This means that a redundant pseudowire pair is configured on both the left and right sides of the VLL cross-connect. A VLL path can be switched to a secondary standby pseudowire at the endpoint node (T-PE) as well as at switching node (S-PE). This requires that a mechanism be built-in to avoid creating loops in the path of the VLL service. This mechanism could be based on the use of the new active/standby status bits and distributing them among the T-PE and S-PE nodes to indicate which spoke SDP to activate for the end-to-end VLL path.

Pseudowire Redundancy for a VPLS Service

There are a number of scenarios which require that a set of redundant pseudowires originate or terminate in a VPLS service at a 7x50 node. An example is when in Figure 26 the service configured in T-PE1 and T-PE2 is a VPLS service while the service configured in the S-PE nodes is an Ethernet VLL with pseudowire switching.

Pseudowire Redundancy Service Models

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

The redundant VLL service model is described in the following section, Redundant VLL Service Model. The selection rules for the active transmit endpoint object are detailed in Section VLL Endpoint Active Transmit Object Selection Rules.

Redundant VLL Service Model

In order to implement pseudowire redundancy, a VLL service accommodates more than a single object on the SAP side and on the spoke SDP side. Figure 27 illustrates the model for a redundant VLL service based on the concept of endpoints.

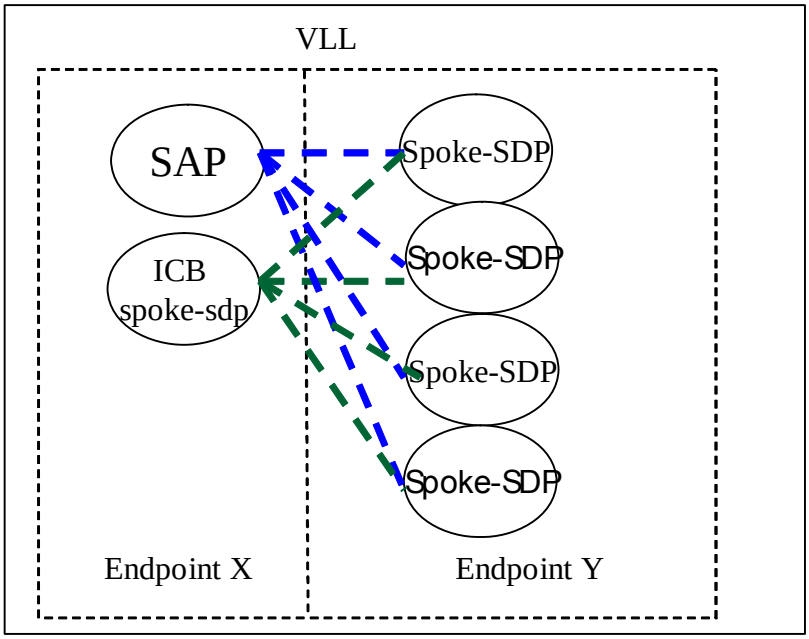

Figure 14: Redundant VLL Endpoint Objects

A VLL service supports by default two implicit endpoints managed internally by the system. Each endpoint can only have one object, a SAP or a spoke SDP.

In order to add more objects, up to two (2) explicitly named endpoints may be created per VLL service. The endpoint name is locally significant to the VLL service. Let's refer to them as endpoint 'X' and endpoint 'Y' as illustrated in Figure 27.

Note that Figure 27 is merely an example and that the “Y” endpoint can also have a SAP and/or an ICB spoke SDP. The following details the four types of endpoint objects supported and the rules used when associating them with an endpoint of a VLL service:

- SAP — There can only be a maximum of one SAP per VLL endpoint.
- Primary spoke SDP — The VLL service always uses this pseudowire and only switches to a secondary pseudowire when it is down the VLL service switches the path to the primary pseudowire when it is back up. The user can configure a timer to delay reverting back to primary or to never revert. There can only be a maximum of one primary spoke SDP per VLL endpoint.
- Secondary spoke SDP — There can be a maximum of four secondary spoke SDP per endpoint. The user can configure the precedence of a secondary pseudowire to indicate the order in which a secondary pseudowire is activated.
- Inter-Chassis Backup (ICB) spoke SDP — Special pseudowire used for MC-LAG and pseudowire redundancy application. Forwarding between ICBs is blocked on the same node. The user has to explicitly indicate the spoke SDP is actually an ICB at creation time. There are however a few scenarios below where the user can configure the spoke SDP as ICB or as a regular spoke SDP on a given node. The CLI for those cases will indicate both options.

A VLL service endpoint can only use a single active object to transmit at any given time but can receive from all endpoint objects

An explicitly named endpoint can have a maximum of one SAP and one ICB. Once a SAP is added to the endpoint, only one more object of type ICB spoke SDP is allowed. The ICB spoke SDP cannot be added to the endpoint if the SAP is not part of a MC-LAG instance. Conversely, a SAP which is not part of a MC-LAG instance cannot be added to an endpoint which already has an ICB spoke SDP.

An explicitly named endpoint, which does not have a SAP object, can have a maximum of four spoke SDPs and can include any of the following:

- A single primary spoke SDP.
- One or many secondary spoke SDPs with precedence.
- A single ICB spoke SDP.

VLL Endpoint Active Transmit Object Selection Rules

The following pseudocode provides the selection rules for the active transmit object on a VLL endpoint. These rules achieve the following desired objectives:

1. ICB on “X” endpoint is a backup for the MC-LAG SAP. VLL path active transmit object is the SAP if it is up and active, or ICB if it is up and did not receive either sdp-binding down status bits or pseudowire not forwarding status bits.
2. ICB in the “Y” endpoint is a backup for active spoke SDP. VLL path active transmit object is the spoke SDP which is up and shows remote active. If active spoke SDP goes locally down or there is no active spoke SDP (for example, all spokes are locally down or showing remote Standby), we will switch to the ICB spoke SDP provided it is up and did not receive sdp-binding down status bits or pseudowire not forwarding status bits. Otherwise, the algorithm selects the best available spoke SDP among those which are not locally down based on the following priority:

   a. Spoke SDP with all received T-LDP status bits clear (0x00000000). This means the remote spoke SDP is active and both remote SAP and spoke SDP are operationally up.

   b. Spoke SDP with received T-LDP status bits indicating remote spoke SDP is in standby (0x00000020) but both remote SAP and spoke SDP are operationally up.

   c. Spoke SDP with received T-LDP status bits indicating remote spoke SDP is active and remote SAP is down (0x00000006).

   d. Spoke SDP with received T-LDP status bits indicating remote spoke SDP is in standby and remote SAP is down (0x00000026).

   e. Spoke SDP with received T-LDP status bits indicating remote spoke SDP is active and in not forwarding (0x00000001) state.

   f. Spoke SDP with received T-LDP status bits indicating remote spoke SDP is in standby and in not forwarding (0x00000021) state.

   g. Spoke SDP with received T-LDP status bits indicating remote spoke SDP is active and is operationally down (0x00000018).

   h. Spoke SDP with received T-LDP status bits indicating remote spoke SDP is in standby and is operationally down (0x00000038).

   i. If multiple objects remain after running steps “a” through “h”, the selected active transmit object is the one with the lowest local precedence value. If multiple objects remain, the selected active transmit object is the one with the lowest spoke SDP identifier

   j. If an active transmit object is found in steps “a” through “i” but its received T-LDP status bits indicate values other than active and operationally up (0x00000000) and SAP down (0x00000006), select ICB spoke SDP. The ICB must be operationally up and did not receive sdp-binding down status bits or pseudowire not forwarding status bits.

   k. If an active transmit object is not found in steps “a” through “i”, select ICB spoke SDP if it is operationally up and did not receive sdp-binding down status bits or pseudowire not forwarding status bits.

3. The only two cases where traffic is discarded locally are when all objects on the endpoint are locally down and when the selection of the active transmit object results in an attempt to forward between two ICB spoke SDPs.

T-LDP Status Notification Handling Rules

Referring to Figure 27 on page 137 as a reference, the following are the rules for generating, processing, and merging T-LDP status notifications in VLL service with endpoints. Note that any allowed combination of objects as specified in Redundant VLL Service Model on page 137 can be used on endpoints “X” and “Y”. The following sections refer to the specific combination objects in Figure 27 as an example to describe the more general rules.

Processing Endpoint SAP Active/Standby Status Bits

The advertised admin forwarding status of active/standby reflects the status of the local LAG SAP in MC-LAG application. If the SAP is not part of a MC-LAG instance, the forwarding status of active is always advertised.

When the SAP in endpoint “X” is part of a MC-LAG instance, a node must send T-LDP forwarding status bit of “SAP active/standby” over all “Y” endpoint spoke SDPs, except the ICB spoke SDP, whenever this status changes. The status bit sent over the ICB is always zero (active by default).

When the SAP in endpoint “X” is not part of a MC-LAG instance, then the forwarding status sent over all “Y” endpoint spoke SDP's should always be set to zero (active by default).

The received SAP active/standby status is saved and used for selecting the active transmit endpoint object as per the pseudo-code in VLL Endpoint Active Transmit Object Selection Rules on page 139.

Processing and Merging Local and Received Endpoint Object Up/Down Operational Status

Endpoint “X” is operationally up if at least one of its objects is operationally up. It is down if all its objects are operationally down.

If the SAP in endpoint “X” transitions locally to the down state, or received a SAP down notification via SAP specific OAM signal, the node must send T-LDP SAP down status bits on the “Y” endpoint ICB spoke SDP only. Note that Ethernet SAP does not support SAP OAM protocol. All other SAP types cannot exist on the same endpoint as an ICB spoke SDP since non Ethernet SAP cannot be part of a MC-LAG instance.

If the ICB spoke SDP in endpoint “X” transitions locally to down state, the node must send T-LDP SDP-binding down status bits on this spoke SDP.

If the ICB spoke SDP in endpoint “X” received T-LDP SDP-binding down status bits or pseudowire not forwarding status bits, the node saves this status and takes no further action. The saved status is used for selecting the active transmit endpoint object.

If all objects in endpoint “X” transition locally to down state, and/or received a SAP down notification via remote T-LDP status bits or via SAP specific OAM signal, and/or received status bits of SDP-binding down, and/or received status bits of pseudowire not forwarding, the node must send status bits of SAP down over all “Y” endpoint spoke SDPs, including the ICB.

Endpoint “Y” is operationally up if at least one of its objects is operationally up. It is down if all its objects are operationally down.

If a spoke SDP in endpoint “Y”, including the ICB spoke SDP, transitions locally to down state, the node must send T-LDP SDP-binding down status bits on this spoke SDP.

If a spoke SDP in endpoint “Y”, including the ICB spoke SDP, received T-LDP SAP down status bits, and/or received T-LDP SDP-binding down status bits, and/or received status bits of pseudowire not forwarding, the node saves this status and takes no further action. The saved status is used for selecting the active transmit endpoint object.

If all objects in endpoint “Y”, except the ICB spoke SDP, transition locally to down state, and/or received T-LDP SAP down status bits, and/or received T-LDP SDP-binding down status bits, and/or received status bits of pseudowire not forwarding, the node must send status bits of SDP-binding down over the “X” endpoint ICB spoke SDP only.

If all objects in endpoint “Y” transition locally to down state, and/or received T-LDP SAP down status bits, and/or received T-LDP SDP-binding down status bits, and/or received status bits of pseudowire not forwarding, the node must send status bits of SDP-binding down over the “X” endpoint ICB spoke SDP, and must send a SAP down notification on the “X” endpoint SAP via the SAP specific OAM signal if applicable. An Ethernet SAP does not support signaling status notifications.

VLL Service Considerations

This section describes various of the general 7750 SR service features and any special capabilities or considerations as they relate to VLL services.

SDPs

The most basic SDPs must have the following:

- A locally unique SDP identification (ID) number.
- The system IP address of the originating and far-end 7750 SR routers.
- An SDP encapsulation type - either GRE or MPLS.

The most basic Apipe and Fpipe SDP configurations must have the following:

- A locally unique SDP identification (ID) number and vc-id.

SDP Statistics for VPLS and VLL services

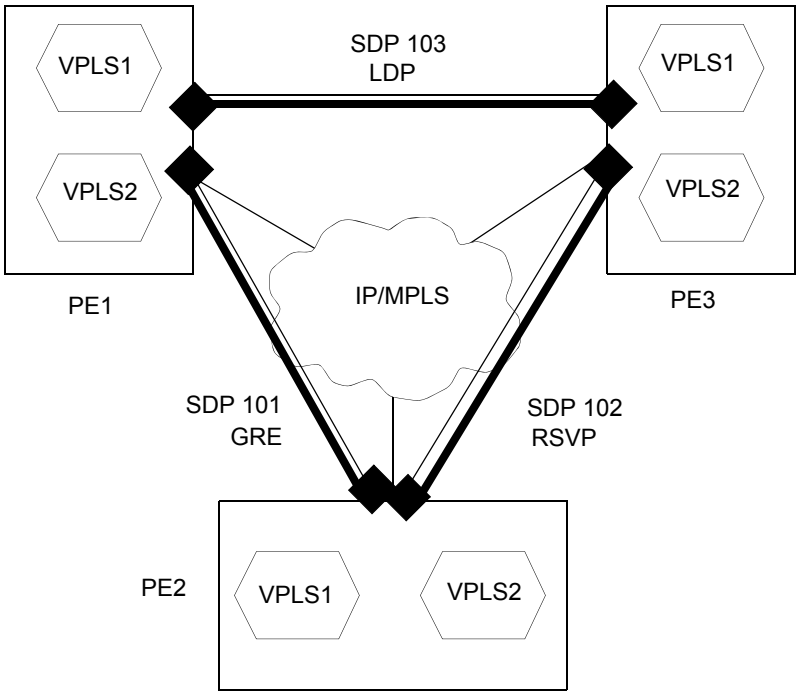

Figure 15: SDP Statistics for VPLS and VLL Services

The simple three-node network described in Figure 28 shows two MPLS SDPs and one GRE SDP defined between the nodes. These SDPs connect VPLS1 and VPLS2 instances that are defined in the three nodes. With this feature the operator will have local CLI based as well as SNMP based statistics collection for each VC used in the SDPs. This will allow for traffic management of tunnel usage by the different services and with aggregation the total tunnel usage.

SAP Encapsulations and Pseudowire Types

The 7750 SR Epipe service is designed to carry Ethernet frame payloads, so it can provide connectivity between any two SAPs that pass Ethernet frames. The following SAP encapsulations are supported on the 7750 SR Epipe service:

- Ethernet null
- Ethernet dot1q
- SONET/SDH BCP-null
- SONET/SDH BCP-dot1q
- ATM VC with RFC 2684 Ethernet-bridged encapsulation (see Ethernet Interworking VLL on page 112)
- FR VC with RFC 2427 Ethernet-bridged encapsulation (see Ethernet Interworking VLL on page 112)

Note that while different encapsulation types can be used, encapsulation mismatching can occur if the encapsulation behavior is not understood by connecting devices and are unable to send and receive the expected traffic. For example if the encapsulation type on one side of the Epipe is dot1q and the other is null, tagged traffic received on the null SAP will be double tagged when it is transmitted out of the dot1q SAP.

ATM VLLs can be configured with both endpoints (SAPs) on the same router or with the two endpoints on different 7750 SRs. In the latter case, Pseudowire Emulation Edge-to-Edge (PWE3) signalling is used to establish a pseudowire between the devices allowing ATM traffic to be tunneled through an MPLS or GRE network:

Two pseudowire encapsulation modes, i.e., SDP vc-type, are available:

- PWE3 N-to-1 Cell Mode Encapsulation
- PWE3 AAL5 SDU Mode Encapsulation

The endpoints of Frame Relay VLLs must be Data-Link Connection Identifiers (DLCIs) on any port that supports Frame Relay. The pseudowire encapsulation, or SDP vc-type, supported is the 1-to-1 Frame Relay encapsulation mode.

PWE3 N-to-1 Cell Mode

The endpoints of an N-to-1 mode VLL can be:

- ATM VCs — VPI/VCI translation is supported (i.e., the VPI/VCI at each endpoint does not need to be the same).
- ATM VPs — VPI translation is supported (i.e., the VPI at each endpoint need not be the same, but the original VCI will be maintained).
- ATM VTs (a VP range) — No VPI translation is supported (i.e., the VPI/VCI of each cell is maintained across the network).
- ATM ports — No translation is supported (i.e., the VPI/VCI of each cell is maintained across the network).

For N-to-1 mode VLLs, cell concatenation is supported. Cells will be packed on ingress to the VLL and unpacked on egress. As cells are being packed, the concatenation process may be terminated by:

- Reaching a maximum number of cells per packet.
- Expiry of a timer.
- (Optionally) change of the CLP bit.
- (Optionally) change of the PTI bits indicating end of AAL5 packet.

In N-to-1 mode, OAM cells are transported through the VLL as any other cell. The PTI and CLP bits are untouched, even if VPI/VCI translation is carried out.

PWE3 AAL5 SDU Mode

The endpoints of an AAL5 SDU mode VLL must be ATM VCs specified by port/vpi/vci. VPI/VCI translation is supported. The endpoint can also be a FR VC, specified by port/dlci. In this case FRF.5 FR-ATM network interworking is performed between the ATM VC SAP or the SDP and the FR VC SAP.

In SDU mode, the mandatory PWE3 control word is supported. This allows the ATM VLL to transport OAM cells along with SDU frames, using the “T” bit to distinguish between them, to recover the original SDU length, and to carry CLP, EFCI and UU information.

QoS Policies

When applied to 7750 SR Epipe, Apipe, and Fpipe services, service ingress QoS policies only create the unicast queues defined in the policy. The multipoint queues are not created on the service.

With Epipe, Apipe, and Fpipe services, 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 a service. QoS policies on Apipes cannot perform any classification and on Fpipes Layer 3 (IP) classification is performed.

Filter Policies

7750 SR Epipe, Fpipe, and Ipipe services can have a single filter policy associated on both ingress and egress. Both MAC and IP filter policies can be used on Epipe services.

Filters cannot be configured on 7750 SR Apipe service SAPs.

MAC Resources

Epipe services are point-to-point layer 2 VPNs capable of carrying any Ethernet payloads. Although an Epipe is a layer 2 service, the 7750 SR Epipe implementation does not perform any MAC learning on the service, so Epipe services do not consume any MAC hardware resources.

Configuring a VLL Service with CLI

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

Topics in this section include:

Ipipe

→ Modifying Ipipe Service Parameters on page 217

→ Disabling an Ipipe Service on page 218

→ Re-enabling an Ipipe Service on page 218

→ Deleting an Ipipe Service on page 219

- Configure endpoint parameters on page 159

List of Commands

Table 4 lists all the service configuration commands indicating the configuration level at which each command is implemented with a short command description. VLL services are configured in the `config>service` context. The command list is organized in the following task-oriented manner:

- Configure an Epipe service
  - → Configure an Epipe SAP
    - – Configure Epipe SAP ATM parameters
    - – Configure Epipe SAP egress parameters
    - – Configure Epipe SAP ingress parameters
  - → Configure an Epipe spoke SDP
- Configure an Apipe service
  - → Configure Apipe service parameters
  - → Configure Apipe SAP parameters
    - – Configure Apipe SAP ATM parameters
  - → Configure Apipe spoke SDP parameters
- Configure an Fpipe service
  - → Configure Fpipe service parameters
  - → Configure Fpipe SAP parameters
  - → Configure Fpipe spoke SDP parameters
- Configure an Ipipe service
  - → Configure an Ipipe SAP
    - – Configure Ipipe SAP ATM parameters
    - – Configure Ipipe SAP egress parameters
    - – Configure Ipipe SAP ingress parameters
  - → Configure Ipipe spoke SDP parameters
- Configure endpoint parameters

Table 1: CLI Commands to Configure VLL Service Parameters

Command	Description	Page
Configure an Epipe service		
`config>service# epipe service-id customer customer-id [vpn vpn-id] [vc-switching]`		
`service-id`	Specifies a unique service identification number identifying the service in the service domain.	235
`customer`	Specifies the customer ID number to be associated with the service.	235
`vpn vpn-id`	Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.	235
`description`	Specifies a text string describing the Epipe service.	234
`service-mtu`	Configures the service payload MTU in bytes for the service ID overriding the service-type default MTU.	239
`no shutdown`	Administratively enables the service.	233
Configure an Epipe SAP		
`config>service>epipe# sap sap-id`		
`accounting-policy`	Associates the accounting policy ID to the SAP. Accounting policies are configured in the `config>log` context.	248
`collect-stats`	Enables the collection of accounting and statistical data for the SAP, network port, or IP interface.	248
`description`	Specifies a text string describing the SAP.	234
`multi-service-site`	Associates a customer site with the *customer-site-name* parameter.	246
`no shutdown`	Administratively enables the SAP.	233
Configure Epipe SAP ATM parameters		
`config>service>epipe>sap# atm`		
`atm`	Enables access to the context to configure ATM-related attributes.	274
`egress`	Configures egress ATM attributes for the SAP.	274
`encapsulation`	Specifies the data encapsulation for an ATM PVCC delimited SAP.	275
`traffic-desc`	Assigns an ATM traffic descriptor profile to a given context (for example a SAP).	275
`ingress`	Configures ingress ATM attributes for the SAP.	274
`oam`	Enables the context to configure OAM functionality for a PVCC delimiting a SAP.	276
`alarm-cells`	Configures AIS/RDI fault management on a PVCC.	276
Configure Epipe SAP egress parameters		
`config>service>epipe>sap# egress`		

Table 1: CLI Commands to Configure VLL Service Parameters (Continued)

Command	Description	Page
filter	Associates an IP filter policy or MAC filter policy with an egress SAP or IP interface.	249
qos	Associates a Quality of Service (QoS) policy with an egress SAP or IP interface.	251
queue-override	Enables the context to configure override values for the specified SAP egress QoS queue.	252
queue	Specifies the ID of the queue whose parameters are to be overridden.	252
avg-frame-overhead	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.	254
adaptation-rule	Specifies attributes of the adaptation rule parameters.	253
cbs	Specifies attributes of the CBS parameters.	255
high-prio-only	Specifies attributes of the high-prio-only parameters.	256
mbs	Specifies attributes of the MBS parameters.	257
rate	Specifies attributes of the Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.	258
scheduler-override	Specifies the set of attributes whose values have been overridden via management on this virtual scheduler.	259
scheduler-policy	Associates an existing scheduler policy to an egress scheduler used by SAP queues associated with this multi-service customer site.	262
Configure Epipe SAP ingress parameters		
config>service>epipe>sap#		
ingress	Configures ingress SAP QoS policies and filter policies.	249
filter	Associates an IP filter policy or MAC filter policy with an ingress SAP or IP interface.	249
match-qinq-dot1p	Configures filtering based on the p-bits in the top or bottom tag of a Q-in-Q encapsulated Ethernet frame.	263
qos	Associates a Quality of Service (QoS) policy with an ingress SAP or IP interface.	251
queue-override	Enables the context to configure override values for the specified SAP egress QoS queue.	252
queue	Specifies the ID of the queue whose parameters are to be overridden.	252
adaptation-rule	Specifies the queue's adaptation rule parameters.	253
cbs	Specifies the queue's CBS parameters.	255
high-prio-only	Specifies the specified high-prio-only parameters.	256
mbs	Specifies the specified MBS parameters.	257

Table 1: CLI Commands to Configure VLL Service Parameters (Continued)

Command	Description	Page
rate	Specifies the specified Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.	258
scheduler-override	Specifies the set of attributes whose values have been overridden via management on this virtual scheduler.	259
scheduler-policy	Associates an existing scheduler policy to an egress scheduler used by SAP queues associated with this multi-service customer site.	259
Configure an Epipe spoke SDP		
config>service>epipe# spoke-sdp *sdp-id:vc-id* [vc-type {ether|vlan}][no-endpoint] config>service>epipe# spoke-sdp *sdp-id:vc-id* [vc-type {ether|vlan}] endpoint *endpoint-name* [icb]		
spoke-sdp	Binds a service to an existing SDP.	125
accounting-policy	Specifies the accounting policy to apply to the SAP.	248
collect-stats	Enables the collection of accounting and statistical data for the SAP or network port.	248
control-word	Specifies whether the use of the control word is preferred or not.	127
ingress	Enables the context to specify the ingress filter policy and VC label value.	129
egress	Enables the context to specify the egress filter policy and VC label value.	129
filter	Associates an IP filter policy with an egress Service Access Point (SAP) or IP interface.	129
vc-label	Configures the egress VC label.	126
vlan-vc-tag	Specifies an explicit dot1q value used encapsulating to the SDP far end.	127
precedence	Specifies the precedence of this SDP bind when there are multiple SDP binds attached to one service endpoint.	128
no shutdown	Administratively enables the SDP.	117
Configure an Apipe service		
config>service# apipe *service-id* [customer *customer-id*] [vpn *vpn-id*] [vc-type {atm-vcc|atm-sdu|atm-vpc|atm-cell}] [vc-switching]		
service-id	Specifies a unique service identification number identifying the service in the service domain.	236
customer	Specifies the existing customer ID number associated with the service.	236
vpn-id	Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.	236

Table 1: CLI Commands to Configure VLL Service Parameters (Continued)

Command	Description	Page
vc-type	Specifies a 15 bit value that defines the type of the VC signaled to the peer.	236
Configure Apipe service parameters		
config>service>apipe#		
description	Specifies a text string describing the service.	234
interworking	Specifies the interworking function that should be applied for packets that ingress/egress SAPs that are part of an Apipe service.	239
service-mtu	Configures the MTU to be used for this service.	239
no shutdown	Administratively enables the Apipe service.	233
Configure Apipe SAP parameters		
config>service>apipe# sap		
accounting-policy	Specifies the accounting policy to apply to the SAP.	248
collect-stats	Enables the collection of accounting and statistical data for the SAP or network port.	248
egress	Enables a context to configure egress SAP Quality of Service (QoS) policies.	249
ingress	Enables a context to configure ingress SAP Quality of Service (QoS) policies.	249
qos	Associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP).	251
queue-override	Enables the context to configure override values for the specified SAP egress QoS queue.	252
queue	Specifies the ID of the queue whose parameters are to be overridden.	252
avg-frame-overhead	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.	254
adaptation-rule	Specifies attributes of the adaptation rule parameters.	253
cbs	Specifies attributes of the CBS parameters.	255
high-prio-only	Specifies attributes of the high-prio-only parameters.	256
mbs	Specifies attributes of the MBS parameters.	257
rate	Specifies attributes of the Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.	258
scheduler-policy	Applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues associated with a multi-service customer site.	262
multi-service-site	Associates a customer site with the *customer-site-name* parameter.	246
no shutdown	Administratively enables the SAP.	233

Table 1: CLI Commands to Configure VLL Service Parameters (Continued)

Command	Description	Page
Configure Apipe SAP ATM parameters		
`config>service>apipe>sap# atm`		
`atm`	Enables access to the context to configure ATM-related attributes.	274
`egress`	Configures egress ATM attributes for the SAP.	274
`traffic-desc`	Assigns an ATM traffic descriptor profile to a given context (for example a SAP).	275
`ingress`	Configures ingress ATM attributes for the SAP.	274
`oam`	Enables the context to configure OAM functionality for a PVCC delimiting a SAP.	276
`alarm-cells`	Configures AIS/RDI fault management on a PVCC.	276
`terminate`	Specifies whether this SAP will act as an OAM termination point. ATM SAPs can be configured to tunnel or terminate OAM cells.	277
Configure Apipe spoke SDP parameters		
`config>service>apipe# spoke-sdp` *`sdp-id:vc-id`* `[no-endpoint]` `config>service>apipe# spoke-sdp` *`sdp-id:vc-id`* `[endpoint` *`endpoint-name`*`] [icb]`		
`cell-concatenation`	Enables the context to provide access to the various options that control the termination of ATM cell concatenation into an MPLS frame. Several options can be configured simultaneously.	269
`aal5-frame-aware`	Enables the AAL5 end-of-message (EOM) to be an indication to complete the cell concatenation operation.	269
`clp-change`	Enables the CLP change to be an indication to complete the cell concatenation operation.	269
`max-cells`	Enables the configuration of the maximum number of ATM cells to accumulate into an MPLS packet.	269
`max-delay`	Enables the configuration of the maximum amount of time to wait while performing ATM cell concatenation into an MPLS packet before transmitting the MPLS packet.	270
`control-word`	Specifies whether the use of the control word is preferred or not.	127
`egress`	Configures the egress spoke SDP context.	271
`ingress`	Configures the ingress spoke SDP context.	271
`precedence`	Specifies the precedence of this SDP bind when there are multiple SDP binds attached to one service endpoint.	128
`vc-label`	Configures the egress or ingress VC label.	272
`no shutdown`	Administratively enables the spoke SDP binding	233
Configure an Fpipe service		

Table 1: CLI Commands to Configure VLL Service Parameters (Continued)

Command	Description	Page
config>service# fpipe *service-id* [customer *customer-id*] [vpn *vpn-id*] [vc-type {fr-dlci}] [vc-switching]		
service-id	Specifies a unique service identification number identifying the service in the service domain.	237
customer	Specifies the existing customer ID number to be associated with the service.	237
vpn-id	Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.	237
vc-type	Specifies a 15 bit value that defines the type of the VC signaled to the peer.	237
Configure Fpipe service parameters		
config>service>fpipe#		
description	Specifies a text string describing the service.	234
endpoint	Specifies a service endpoint.	238
service-mtu	Configures the MTU to be used for this service.	239
no shutdown	Administratively enables the Fpipe service.	233
Configure Fpipe SAP parameters		
config>service>fpipe# sap		
accounting-policy	Specifies the accounting policy to apply to the SAP.	248
collect-stats	Enables the collection of accounting and statistical data for the SAP or network port.	248
egress	Enables a context to configure egress SAP Quality of Service (QoS) policies.	249
ingress	Enables a context to configure ingress SAP Quality of Service (QoS) policies.	249
filter	Associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface.	249
qos	Associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP).	251
queue-override	Enables the context to configure override values for the specified SAP egress QoS queue.	252
queue	Specifies the ID of the queue whose parameters are to be overridden.	252
avg-frame-overhead	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.	254
adaptation-rule	Specifies attributes of the adaptation rule parameters.	253

Table 1: CLI Commands to Configure VLL Service Parameters (Continued)

Command	Description	Page
`cbs`	Specifies attributes of the CBS parameters.	255
`high-prio-only`	Specifies attributes of the high-prio-only parameters.	256
`mbs`	Specifies attributes of the MBS parameters.	257
`rate`	Specifies attributes of the Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.	258
`scheduler-policy`	Applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues associated with a multi-service customer site.	262
`multi-service-site`	Creates a new customer site or edits an existing customer site with the *customer-site-name* parameter.	246
`no shutdown`	Administratively enables the SAP.	233
Configure Fpipe spoke SDP parameters		
`config>service>fpipe# spoke-sdp` *`sdp-id:vc-id`* `[no-endpoint]`		
`config>service>fpipe# spoke-sdp` *`sdp-id:vc-id`* `[endpoint` *`endpoint-name`*`] [icb]`		
`egress`	Configures the egress spoke SDP context.	271
`ingress`	Configures the ingress spoke SDP context.	271
`filter`	Associates an IP filter policy with an ingress or egress Service Distribution Point (SDP).	271
`vc-label`	Configures the egress or ingress VC label.	272
`precedence`	Specifies the precedence of this SDP bind when there are multiple SDP binds attached to one service endpoint.	128
`no shutdown`	Administratively enables the spoke SDP binding	233
Configure an Ipipe service		
`config>service# ipipe` *`service-id`* `[customer` *`customer-id`*`][vpn` *`vpn-id`*`] [vc-switching]`		
`service-id`	Specifies a unique service identification number identifying the service in the service domain.	237
`customer`	Specifies the customer ID number to be associated with the service.	237
`vpn` *`vpn-id`*	Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.	237
`no shutdown`	Administratively enables the service.	233
`description`	Specifies a text string describing the Ipipe service.	233
`endpoint`	Specifies a service endpoint.	238
Configure an Ipipe SAP		
`config>service>ipipe# sap` *`sap-id`*		
`accounting-policy`	Associates the accounting policy ID to the SAP. Accounting policies are configured in the `config>log` context.	248

Table 1: CLI Commands to Configure VLL Service Parameters (Continued)

Command	Description	Page
ce-address	Specifies the IP address of the CE device associated with an Ipipe SAP.	276
collect-stats	Enables the collection of accounting and statistical data for the SAP, network port, or IP interface.	248
description	Specifies a text string describing the SAP.	234
mac	Assigns a specific MAC address to an Ipipe SAP.	247
mac-refresh	Specifies the interval between ARP requests sent on this Ipipe SAP.	247
multi-service-site	Associates the SAP with a customer site name.	246
no shutdown	Administratively enables the SAP.	233
Configure Ipipe SAP ATM parameters		
config>service>ipipe>sap# atm		
atm	Enables access to the context to configure ATM-related attributes.	274
egress	Configures egress ATM attributes for the SAP.	274
traffic-desc	Assigns an ATM traffic descriptor profile to a given context (for example a SAP).	275
ingress	Configures ingress ATM attributes for the SAP.	274
oam	Enables the context to configure OAM functionality for a PVCC delimiting a SAP.	276
alarm-cells	Configures AIS/RDI fault management on a PVCC.	276
encapsulation	Specifies the data encapsulation for an ATM PVCC delimited SAP.	275
Configure Ipipe SAP egress parameters		
config>service>ipipe>sap#		
egress	Configures egress SAP QoS policies and filter policies.	249
filter	Associates an IP filter policy or MAC filter policy with an egress SAP or IP interface.	249
qos	Associates a Quality of Service (QoS) policy with an egress SAP or IP interface.	251
queue-override	Enables the context to configure override values for the specified SAP egress QoS queue.	252
queue	Specifies the ID of the queue whose parameters are to be overridden.	252
adaptation-rule	Specifies attributes of the adaptation rule parameters.	253
cbs	Specifies attributes of the CBS parameters.	255
high-prio-only	Specifies attributes of the high-prio-only parameters.	256
mbs	Specifies attributes of the MBS parameters.	257

Table 1: CLI Commands to Configure VLL Service Parameters (Continued)

Command	Description	Page
rate	Specifies attributes of the Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.	258
scheduler-override	Specifies attributes of the the set of attributes whose values have been overridden via management on this virtual scheduler.	259
scheduler-policy	Associates an existing scheduler policy to an egress scheduler used by SAP queues associated with this multi-service customer site.	262
Configure Ipipe SAP ingress parameters		
config>service>ipipe>sap#		
ingress	Configures ingress SAP QoS policies and filter policies.	249
filter	Associates an IP filter policy or MAC filter policy with an ingress SAP or IP interface.	249
match-qinq-dot1p	Configures filtering based on the p-bits in the top or bottom tag of a Q-in-Q encapsulated Ethernet frame.	263
qos	Associates a Quality of Service (QoS) policy with an ingress SAP or IP interface.	251
queue-override	Enables the context to configure override values for the specified SAP egress QoS queue.	252
queue	Specifies the ID of the queue whose parameters are to be overridden.	252
adaptation-rule	Specifies attributes of the queue's adaptation rule parameters.	253
cbs	Specifies attributes of the CBS parameters.	255
high-prio-only	Specifies attributes of the high-prio-only parameters.	256
mbs	Specifies attributes of the MBS parameters.	257
rate	Specifies attributes of the Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.	258
scheduler-override	Specifies the set of attributes whose values have been overridden via management on this virtual scheduler.	259
scheduler-policy	Associates an existing scheduler policy to an egress scheduler used by SAP queues associated with this multi-service customer site.	259
Configure Ipipe spoke SDP parameters		
config>service>ipipe# spoke-sdp *sdp-id:vc-id* [no-endpoint]		
config>service>ipipe# spoke-sdp *sdp-id:vc-id* [endpoint *endpoint-name*] [icb]		
ce-address	Specifies the IP address of the CE device associated with an Ipipe spoke SDP.	276
control-word	Specifies whether the use of the control word is preferred or not.	127
egress	Configures the egress spoke SDP context.	271
ingress	Configures the ingress spoke SDP context.	249

Table 1: CLI Commands to Configure VLL Service Parameters (Continued)

Command	Description	Page
`filter`	Associates an IP filter policy with an ingress or egress Service Distribution Point (SDP).	249
`vc-label`	Configures the egress or ingress VC label.	272
`precedence`	Specifies the precedence of this SDP bind when there are multiple SDP binds attached to one service endpoint.	128
`no shutdown`	Administratively enables the spoke SDP binding	233
Configure endpoint parameters		

`config>service>apipe`
`config>service>epipe`
`config>service>fpipe`
`config>service>ipipe`		

`endpoint`	Specifies a service endpoint.	238
`active-hold-delay`	Specifies that the node will delay sending the change in the T-LDP status bits for the VLL endpoint when the MC-LAG transitions the LAG subgroup which hosts the SAP for this VLL endpoint from "active" to "standby" or when any object in the endpoint.	238
`description`	Specifies a text string describing the endpoint.	233
`revert-time`	Configures the time to wait before reverting back to the primary spoke SDP defined on this service endpoint, after having failed over to a backup spoke SDP.	239

Basic Configurations

Configuring an Epipe Service

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

- Customer ID (see Configuring Customers on page 64)
- For a local service, two SAPs must be configured specifying the source and destination ports
- For a distributed service, one SAP and one SDP must be specified

The following example displays a sample configuration of a local Epipe service configured on ALA-12.

```
*A:ALA-12>config>service# info
#------------------------------------------
...
        epipe 7 customer 6 vpn 7 create
            description "Local epipe service"
            sap 1/1/6:0 create
            exit
            sap 2/1/8:0 create
            exit
        exit
...
#------------------------------------------
*A:ALA-12>config>service#
```

The following example displays a sample configuration of a distributed Epipe service between ALA-1 and ALA-2.

```
*A:ALA-1>config>service>epipe# info
----------------------------------------------
            description "Distributed Epipe service to east coast"
            sap 5/1/2:0 create
```

```
            exit
            spoke-sdp 2:6 create
                ingress
                    vc-label 6300
                exit
                egress
                    vc-label 6298
                exit
            exit
            no shutdown
-----------------------------------------------
*A:ALA-1>config>service>epipe#


*A:ALA-2>config>service>epipe# info
-----------------------------------------------
            description "Distributed Epipe service to west coast"
            sap 5/1/2:0 create
            exit
            spoke-sdp 2:6 create
                ingress
                    vc-label 6298
                exit
                egress
                    vc-label 6300
                exit
            exit
            no shutdown
-----------------------------------------------
*A:ALA-2>config>service>epipe#
```

Configuring an Apipe Service

The following fields require specific input (there are no defaults) to configure basic Apipe services:

- Customer ID (refer to Configuring Customers on page 64)
- Specify SAP parameters
- Specify spoke SDP parameters

The following example displays sample configurations of Apipe services.

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
----------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 1/1/1:0/32 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:4 create
            exit
            no shutdown
        exit
...
----------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
----------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 2/2/2:0/32 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:4 create
            exit
            no shutdown
        exit
...
```

Configuring an Fpipe Service

The following fields require specific input (there are no defaults) to configure basic Fpipe services:

- Customer ID (refer to Configuring Customers on page 64)
- Specify SAP parameters
- Specify spoke SDP parameters

The following example displays sample configurations of Fpipe services.

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
----------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 1/2/1:16 create
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
            exit
            no shutdown
        exit
...
----------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
----------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 2/1/1.1:16 create
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
            exit
            no shutdown
        exit
```

```
...
-------------------------------------
A:ALA-42>config>service#
```

Configuring an Ipipe Service

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

- Customer ID (refer to Configuring Customers on page 64)
- Specify SAP parameters
- Specify spoke SDP parameters

The following example displays sample configurations of Ipipe services.

```
A:ALA-280>config>service>ipipe# info
----------------------------------------------
            sap 1/1/4:0 create
                ingress
                    qos 2
                exit
                egress
                    qos 1010
                exit
                ce-address 10.40.30.10
                mac 14:30:01:01:00:04
            exit
            spoke-sdp 3:100 create
                ingress
                    filter ip 10
                exit
                ce-address 10.40.30.11
            exit
            no shutdown
----------------------------------------------
A:ALA-280>config>service>ipipe#


A:ALA-49>config>service>ipipe# info
----------------------------------------------
            sap 2/1/4:0 create
                ingress
                    qos 2
                exit
                egress
                    qos 1010
                exit
                ce-address 128.251.10.10
            exit
            spoke-sdp 8:200 create
                ingress
                    filter ip 10
                exit
                ce-address 128.251.10.11
            exit
            no shutdown
----------------------------------------------
A:ALA-49>config>service>ipipe#
```

Common Configuration Tasks

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

- Associate the service with a customer ID.
- Define SAP parameters
  - → Optional - configure ATM parameters
  - → Optional - select egress and ingress QoS and/or scheduler policies (configured in `config>qos` context).
  - → Optional - select accounting policy (configured in `config>log` context).
- Define spoke SDP parameters.
- Enable the service.

Configuring VLL Components

This section provides VLL configuration examples for the VLL services:

Creating an Epipe Service

Use the following CLI syntax to create an Epipe service.

CLI Syntax: config>service# epipe *service-id* [customer *customer-id*] [vpn *vpn-id*] [vc-switching]
description *description-string*
no shutdown

The following example displays the command usage to create an Epipe service:

Example:config>service# epipe **500** customer **5** create
config>service>epipe$ description "**Local epipe service**"
config>service>epipe# no shutdown

The following example displays the Epipe configuration:

```
ALA-1>config>service# info
----------------------------------------
...
        epipe 500 customer 5 vpn 500 create
            description "Local epipe service"
            no shutdown
        exit
----------------------------------------
ALA-1>config>service#
```

Configuring Epipe SAP Parameters

A default QoS policy is applied to each ingress and egress SAP. Additional QoS policies can be configured in the `config>qos` context. Filter policies are configured in the `config>filter` context and explicitly applied to a SAP. There are no default filter policies.

Use the following CLI syntax to create:

- Local Epipe SAPs on page 168
- Distributed Epipe SAPs on page 170

CLI Syntax:
```
config>service# epipe service-id [customer customer-id]
        sap sap-id [endpoint endpoint-name]
        sap sap-id [no-endpoint]
              accounting-policy policy-id
              collect-stats
              description description-string
              no shutdown
              egress
                 filter {ip ip-filter-name|mac mac-filter-name}
                 qos sap-egress-policy-id
                 scheduler-policy scheduler-policy-name
              ingress
                 filter {ip ip-filter-name|mac mac-filter-name}
                 match-qinq-dot1p {top|bottom}
                 qos sap-egress-policy-id
                 scheduler-policy scheduler-policy-name
```

Local Epipe SAPs

To configure a basic local Epipe service, enter the `sap sap-id` command twice with different port IDs in the same service configuration.

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

An existing scheduler policy can be applied to ingress and egress SAPs to be used by the SAP queues. The schedulers comprising the policy are created at the time the scheduler policy is applied to the SAP. If any orphaned queues (queues with a non-existent local scheduler defined) exist on a SAP and the policy application creates the required scheduler, the status on the queue becomes non-orphaned at this time.

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

This example displays the SAP configurations for local Epipe service 500 on SAP 1/1/2 and SAP 1/1/3 on ALA-1.

```
ALA-1>config>service# epipe 500 customer 5 create
     config>service>epipe$ description "Local epipe service"
     config>service>epipe# sap 1/1/2:0 create
     config>service>epipe>sap? ingress
     config>service>epipe>sap>ingress# qos 20
     config>service>epipe>sap>ingress# filter ip 1
     config>service>epipe>sap>ingress# exit
     config>service>epipe>sap# egress
     config>service>epipe>sap>egress# qos 20
     config>service>epipe>sap>egress# scheduler-policy test1
     config>service>epipe>sap>egress# exit
     config>service>epipe>sap# no shutdown
     config>service>epipe>sap# exit

     config>service>epipe# sap 1/1/3:0 create
     config>service>epipe>sap# ingress
     config>service>epipe>sap>ingress# qos 555
     config>service>epipe>sap>ingress# filter ip 1
     config>service>epipe>sap>ingress# exit
     config>service>epipe>sap# egress
     config>service>epipe>sap>egress# qos 627
     config>service>epipe>sap>egress# scheduler-policy alpha
     config>service>epipe>sap>egress# exit
     config>service>epipe>sap# no shutdown
     config>service>epipe>sap# exit
```

The following example displays the local Epipe configuration:

```
ALA-1>config>service# info
----------------------------------------------
...
        epipe 500 customer 5 vpn 500 create
            description "Local epipe service"
            sap 1/1/2:0 create
                ingress
                    qos 20
                    filter ip 1
                exit
                egress
                    scheduler-policy "test1"
                    qos 20
                exit
            exit
            sap 1/1/3:0 create
                ingress
                    qos 555
                    filter ip 1
                exit
                egress
                    scheduler-policy "alpha"
                    qos 627
                exit
            exit
            no shutdown
        exit
----------------------------------------------
ALA-1>config>service#
```

Distributed Epipe SAPs

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

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

An existing scheduler policy can be applied to ingress and egress SAPs to be used by the SAP queues. The schedulers comprising the policy are created at the time the scheduler policy is applied to the SAP. If any orphaned queues (queues with a non-existent local scheduler defined) exist on a SAP and the policy application creates the required scheduler, the status on the queue becomes non-orphaned at this time.

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

For SDP configuration information, see Configuring an SDP on page 68. For SDP binding information, see Configuring SDP Bindings on page 173.

This example configures a distributed service between ALA-1 and ALA-2.

```
ALA-1>epipe 5500 customer 5 create
     config>service>epipe$ description "Distributed epipe service to east coast"
     config>service>epipe# sap 2/1/3:21 create
     config>service>epipe>sap# ingress
     config>service>epipe>sap>ingress# qos 555
     config>service>epipe>sap>ingress# filter ip 1
     config>service>epipe>sap>ingress# exit
     config>service>epipe>sap# egress
     config>service>epipe>sap>egress# qos 627
     config>service>epipe>sap>egress# scheduler-policy alpha
     config>service>epipe>sap>egress# exit
     config>service>epipe>sap# no shutdown
     config>service>epipe>sap# exit
     config>service>epipe#

ALA-2>config>service# epipe 5500 customer 5 create
     config>service>epipe$ description "Distributed epipe service to west coast"
     config>service>epipe# sap 4/1/4:550 create
     config>service>epipe>sap# ingress
     config>service>epipe>sap>ingress# qos 654
     config>service>epipe>sap>ingress# filter ip 1020
     config>service>epipe>sap>ingress# exit
     config>service>epipe>sap# egress
     config>service>epipe>sap>egress# qos 432
     config>service>epipe>sap>egress# filter ip 6
     config>service>epipe>sap>egress# scheduler-policy test1
     config>service>epipe>sap>egress# exit
     config>service>epipe>sap# no shutdown
     config>service>epipe#
```

The following example displays the SAP configurations for ALA-1 and ALA-2:

```
ALA-1>config>service# info
----------------------------------------------
...
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to east coast"
            sap 2/1/3:21 create
                ingress
                    qos 555
                    filter ip 1
                exit
                egress
                    scheduler-policy "alpha"
                    qos 627
                exit
            exit
        exit
...
----------------------------------------------
ALA-1>config>service#


ALA-2>config>service# info
----------------------------------------------
...
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to west coast"
            sap 4/1/4:550 create
                ingress
                    qos 654
                    filter ip 1020
                exit
                egress
                    scheduler-policy "test1"
                    qos 432
                    filter ip 6
                exit
            exit
        exit
...
----------------------------------------------
ALA-2>config>service#
```

Configuring Ingress and Egress SAP Parameters

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

An existing scheduler policy can be applied to ingress and egress SAPs to be used by the SAP queues. The schedulers comprising the policy are created at the time the scheduler policy is applied to the SAP. If any orphaned queues (queues with a non-existent local scheduler defined) exist on a SAP and the policy application creates the required scheduler, the status on the queue becomes non-orphaned at this time.

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

This example displays SAP ingress and egress parameters.

```
ALA-1>config>service# epipe 5500
     config>service>epipe# sap 2/1/3:21
     config>service>epipe>sap# ingress
     config>service>epipe>sap>ingress# qos 555
     config>service>epipe>sap>ingress# filter ip 1
     config>service>epipe>sap>ingress# exit
     config>service>epipe>sap# egress
     config>service>epipe>sap>egress# qos 627
     config>service>epipe>sap>egress# scheduler-policy alpha
     config>service>epipe>sap>egress# exit
     config>service>epipe>sap#
```

The following example displays the Epipe SAP ingress and egress configuration:

```
ALA-1>config>service#
-----------------------------------------------
...
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to east coast"
            sap 2/1/3:21 create
                ingress
                    qos 555
                    filter ip 1
                exit
                egress
                    scheduler-policy "alpha"
                    qos 627
                exit
            exit
            spoke-sdp 2:123 create
                ingress
                    vc-label 6600
                exit
                egress
                    vc-label 5500
                exit
            exit
            no shutdown
        exit
-----------------------------------------------
ALA-1>config>service#
```

Configuring SDP Bindings

Figure 29 displays an example of a distributed Epipe service configuration between two routers, identifying the service and customer IDs, and the uni-directional SDPs required to communicate to the far-end routers.

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

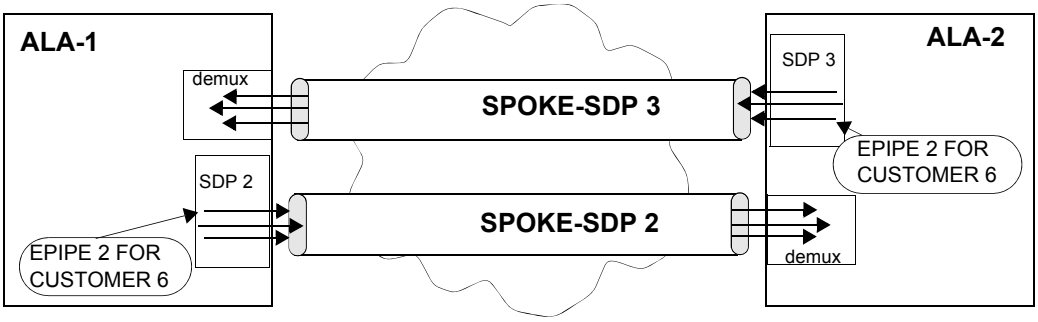

Figure 1: SDPs — Uni-Directional Tunnels

Use the following CLI syntax to create a spoke SDP binding with an Epipe service:

CLI Syntax:
```
config>service# epipe service-id [customer customer-id]
          spoke-sdp sdp-id:vc-id [vc-type {ether|vlan}]
             vlan-vc-tag 0..4094
             egress
                filter {ip ip-filter-id}
                vc-label egress-vc-label
             ingress
                filter {ip ip-filter-id}
                vc-label ingress-vc-label
             no shutdown
```

The following example displays the command usage to bind an Epipe service between ALA-1 and ALA-2. This example assumes the SAPs have already been configured (see Distributed Epipe SAPs on page 170).

```
ALA-1>config>service# epipe 5500
     config>service>epipe# spoke-sdp 2:123
     config>service>epipe>spoke-sdp# egress
     config>service>epipe>spoke-sdp>egress# vc-label 5500
     config>service>epipe>spoke-sdp>egress# exit
     config>service>epipe>spoke-sdp# ingress
```

```
config>service>epipe>spoke-sdp>ingress# vc-label 6600
config>service>epipe>spoke-sdp>ingress# exit
config>service>epipe>spoke-sdp# no shutdown

ALA-2>config>service# epipe 5500
config>service>epipe# spoke-sdp 2:456
config>service>epipe>spoke-sdp# egress
config>service>epipe>spoke-sdp>egress# vc-label 6600
config>service>epipe>spoke-sdp>egress# exit
config>service>epipe>spoke-sdp# ingress
config>service>epipe>spoke-sdp>ingress# vc-label 5500
config>service>epipe>spoke-sdp>ingress# exit
config>service>epipe>spoke-sdp# no shutdown
```

This example displays the SDP binding for the Epipe service between ALA-1 and ALA-2:

```
ALA-1>config>service# info
----------------------------------------------
...
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to east coast"
            sap 2/1/3:21 create
                ingress
                    qos 555
                    filter ip 1
                exit
                egress
                    scheduler-policy "alpha"
                    qos 627
                exit
            exit
            spoke-sdp 2:123 create
                ingress
                    vc-label 6600
                exit
                egress
                    vc-label 5500
                exit
            exit
            no shutdown
        exit
...
----------------------------------------------
ALA-1>config>service#


ALA-2>config>service# info
----------------------------------------------
...
exit
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to west coast"
            sap 4/1/4:550 create
                ingress
                    qos 654
                    filter ip 1020
                exit
                egress
                    scheduler-policy "test1"
                    qos 432
                    filter ip 6
```

```
                exit
            exit
            spoke-sdp 2:456 create
                ingress
                    vc-label 5500
                exit
                egress
                    vc-label 6600
                exit
            exit
            no shutdown
        exit
...
-----------------------------------------------
ALA-2>config>service#
```

Creating an Apipe Service

Use the following CLI syntax to create an Apipe service.

CLI Syntax: config>service# apipe *service-id* [customer *customer-id*] [vpn *vpn-id*] [vc-type {atm-vcc|atm-sdu|atm-vpc|atm-cell}][vc-switching]
description *description-string*
interworking {frf-5}
service-mtu *octets*
no shutdown

The following example displays the command usage to create an Apipe service:

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# apipe 5 customer 1 create
A:ALA-41config>service>apipe# description "apipe test"
A:ALA-41config>service>apipe# service-mtu 1400
A:ALA-41config>service>apipe# no shutdown
A:ALA-41config>service>apipe#
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# apipe 5 customer 1 create
A:ALA-42>config>service>apipe# description "apipe test"
A:ALA-42>config>service>apipe# service-mtu 1400
A:ALA-42>config>service>apipe# no shutdown
A:ALA-42>config>service>apipe#
```

The following example displays the Apipe service creation output.

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            no shutdown
        exit
...
------------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            no shutdown
        exit
...
------------------------------------
A:ALA-42>config>service#
```

Configuring Apipe SAP Parameters

Use the following CLI syntax to configure Apipe SAP parameters.

CLI Syntax: config>service# apipe *service-id* [customer *customer-id*] [vpn *vpn-id*] [vc-type {atm-vcc|atm-sdu|atm-vpc|atm-cell}][vc-switching]
sap *sap-id*
accounting-policy *acct-policy-id*
atm
egress
traffic-desc *traffic-desc-profile-id*
ingress
traffic-desc *traffic-desc-profile-id*
oam
alarm-cells
terminate
collect-stats
description *description-string*
egress
qos *policy-id*
scheduler-policy *scheduler-policy-name*
ingress
qos *policy-id* [shared-queuing]
scheduler-policy *scheduler-policy-name*
multi-service-site *customer-site-name*
no shutdown

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

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# sap 1/1/1:0/32 create
A:ALA-41>config>service>apipe>sap# ingress
A:ALA-41>config>service>apipe>sap>ingress# qos 102
A:ALA-41>config>service>apipe>sap>ingress# exit
A:ALA-41>config>service>apipe>sap# egress
A:ALA-41>config>service>apipe>sap>egress# qos 103
A:ALA-41>config>service>apipe>sap>egress# exit
A:ALA-41>config>service>apipe>sap# no shutdown
A:ALA-41>config>service>apipe>sap# exit
A:ALA-41>config>service>apipe#
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# apipe 5
A:ALA-42>config>service>apipe# sap 2/2/2:0/32 create
A:ALA-42>config>service>apipe>sap# ingress
A:ALA-42>config>service>apipe>sap>ingress# qos 102
A:ALA-42>config>service>apipe>sap>ingress# exit
A:ALA-42>config>service>apipe>sap# egress
```

```
A:ALA-42>config>service>apipe>sap>egress# qos 103
A:ALA-42>config>service>apipe>sap>egress# exit
A:ALA-42>config>service>apipe>sap# no shutdown
A:ALA-42>config>service>apipe>sap# exit
A:ALA-42>config>service>apipe#
```

The following output displays the Apipe SAP configuration.

PE Router 1 (ALA-41):

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

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 2/2/2:0/32 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-42>config>service#
```

Configuring Apipe SDP Bindings

Use the following CLI syntax to create a spoke SDP binding with an Apipe service:

CLI Syntax: config>service# apipe *service-id* [customer *customer-id*] [vpn *vpn-id*] [vc-type {atm-vcc|atm-sdu|atm-vpc|atm-cell}] [vc-switching]
spoke-sdp *sdp-id*:*vc-id*
cell-concatenation
aal5-frame-aware
clp-change
max-cells *cell-count*
max-delay *delay-time*
egress
vc-label *egress-vc-label*
ingress
vc-label *ingress-vc-label*
no shutdown

The following example displays the command usage to create Apipe spoke SDPs:

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# spoke-sdp 1:5 create
A:ALA-41>config>service>apipe>spoke-sdp# no shutdown
A:ALA-41>config>service>apipe>spoke-sdp# exit
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# apipe 5
A:ALA-42>config>service>apipe# spoke-sdp 1:5 create
A:ALA-42>config>service>apipe>spoke-sdp# no shutdown
A:ALA-42>config>service>apipe>spoke-sdp# exit
```

The following output displays the Apipe spoke SDP configurations.

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
----------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 1/1/1:0/32 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
```

```
            spoke-sdp 1:5 create
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 2/2/2:0/32 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:5 create
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-42>config>service#
```

Creating an Fpipe Service

Use the following CLI syntax to create an Fpipe service.

CLI Syntax: config>service# fpipe *service-id* [customer *customer-id*] [vpn *vpn-id*] [vc-type {fr-dlci}][vc-switching]
description *description-string*
service-mtu *octets*
no shutdown

The following example displays the command usage to create an Fpipe service:

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# fpipe 1 customer 1 create
A:ALA-41config>service>fpipe# description "fpipe test"
A:ALA-41config>service>fpipe# service-mtu 1400
A:ALA-41config>service>fpipe# no shutdown
A:ALA-41config>service>fpipe#
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# fpipe 1 customer 1 create
A:ALA-42>config>service>fpipe# description "fpipe test"
A:ALA-42>config>service>fpipe# service-mtu 1400
A:ALA-42>config>service>fpipe# no shutdown
A:ALA-42>config>service>fpipe#
```

The following example displays the Fpipe service creation output.

PE router 1 (A:ALA-41):
```
A:ALA-41>config>service# info
----------------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            no shutdown
        exit
...
----------------------------------------
A:ALA-41>config>service#
```

PE router 2 (A:ALA-42):
```
A:ALA-42>config>service# info
----------------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            no shutdown
        exit
...
----------------------------------------
A:ALA-42>config>service#
```

Configuring Fpipe SAP Parameters

Use the following CLI syntax to configure Fpipe SAP parameters.

CLI Syntax: config>service# fpipe *service-id* [customer *customer-id*] [vpn *vpn-id*] [vc-type {fr-dlci}] [vc-switching]
sap *sap-id*
accounting-policy *acct-policy-id*
collect-stats
description *description-string*
egress
filter [ip *ip-filter-id*]
qos *policy-id*
scheduler-policy *scheduler-policy-name*
ingress
filter [ip *ip-filteri-d*]
qos *policy-id* [shared-queuing]
scheduler-policy *scheduler-policy-name*
multi-service-site *customer-site-name*
no shutdown

The following example displays the command usage to create an Fpipe SAP:

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# sap 1/2/1:16 create
A:ALA-41>config>service>fpipe>sap# ingress
A:ALA-41>config>service>fpipe>sap>ingress# qos 101
A:ALA-41>config>service>fpipe>sap>ingress# exit
A:ALA-41>config>service>fpipe>sap# egress
A:ALA-41>config>service>fpipe>sap>egress# qos 1020
A:ALA-41>config>service>fpipe>sap>egress# exit
A:ALA-41>config>service>fpipe>sap# no shutdown
A:ALA-41>config>service>fpipe>sap# exit
A:ALA-41>config>service>fpipe#
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# fpipe 1
A:ALA-42>config>service>fpipe# sap 2/1/1.1:16 create
A:ALA-42>config>service>fpipe>sap# ingress
A:ALA-42>config>service>fpipe>sap>ingress# qos 101
A:ALA-42>config>service>fpipe>sap>ingress# exit
A:ALA-42>config>service>fpipe>sap# egress
A:ALA-42>config>service>fpipe>sap>egress# qos 1020
A:ALA-42>config>service>fpipe>sap>egress# exit
A:ALA-42>config>service>fpipe>sap# no shutdown
A:ALA-42>config>service>fpipe>sap# exit
A:ALA-42>config>service>fpipe#
```

The following example displays the Fpipe SAP configurations.

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
------------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 1/2/1:16 create
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
------------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 2/1/1.1:16 create
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-42>config>service#
```

Configuring Fpipe SDP Bindings

Use the following CLI syntax to create a spoke SDP binding with an Fpipe service:

CLI Syntax: config>service# fpipe *service-id* [customer *customer-id*] [vpn *vpn-id*] [vc-type {fr-dlci}][vc-switching]
spoke-sdp *sdp-id*:*vc-id*
egress
filter ip *ip-filter-id*
vc-label *egress-vc-label*
ingress
filter ip *ip-filter-id*
vc-label *ingress-vc-label*
no shutdown

The following example displays the command usage to create an Fpipe spoke SDP:

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# spoke-sdp 1:1 create
A:ALA-41>config>service>spoke-sdp# no shutdown
A:ALA-41>config>service>spoke-sdp# exit
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# fpipe 1
A:ALA-42>config>service>fpipe# spoke-sdp 1:1 create
A:ALA-42>config>service>spoke-sdp# no shutdown
A:ALA-42>config>service>spoke-sdp# exit
```

The following output displays the Fpipe spoke SDP configuration.

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
----------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 1/2/1:16 create
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
            exit
            no shutdown
        exit
...
----------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
----------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 2/1/1.1:16 create
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
            exit
            no shutdown
        exit
...
----------------------------------
A:ALA-42>config>service#
```

Creating an Ipipe Service

Use the following CLI syntax to create an Ipipe service.

CLI Syntax: config>service# ipipe *service-id* [customer *customer-id*] [vpn *vpn-id*][vc-switching]
description *description-string*
no shutdown

The following example displays the command usage to create an Epipe service:

Example:
```
config>service# ipipe 202 customer 1 create
    config>service>ipipe$ description "eth_ipipe"
    config>service>ipipe# no shutdown
```

The following example displays the Epipe configuration:

```
ALA-1>config>service# info
----------------------------------------
...
        ipipe 202 customer 1 create
            description "eth_ipipe"
            no shutdown
        exit
----------------------------------------
ALA-1>config>service#
```

Configuring Ipipe SAP Parameters

Configuring ATM to Ethernet local Ipipe example:

Example:
```
config>service# ipipe 202 customer 1 create
config>service>ipipe$ sap 1/1/2:444 create
config>service>ipipe>sap$ description "eth_ipipe"
config>service>ipipe>sap$ ce-address 31.31.31.1
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# sap 1/3/2:445 create
config>service>ipipe>sap$ description "eth_ipipe"
config>service>ipipe>sap$ ce-address 31.31.31.2
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
```

```
A:ALA-48>config>service# info
----------------------------------------------
...
        ipipe 202 customer 1 create
            sap 1/1/2:444 create
                description "eth_ipipe"
                ce-address 31.31.31.1
            exit
            sap 1/3/2:445 create
                description "eth_ipipe"
                ce-address 31.31.31.2
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-48>config>service#
```

Configuring Frame Relay to Ethernet local Ipipe example:

Example:
```
config>service# ipipe 204 customer 1 create
config>service>ipipe$ sap 1/1/2:446 create
config>service>ipipe>sap$ description "eth_fr_ipipe"
config>service>ipipe>sap$ ce-address 32.32.32.1
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# sap 2/2/2:16 create
config>service>ipipe>sap$ ce-address 32.32.32.2
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# no shutdown
config>service>ipipe# exit
config>service#
```

```
A:ALA-48>config>service# info
----------------------------------------------
...
        ipipe 204 customer 1 create
            sap 1/1/2:446 create
                description "eth_fr_ipipe"
                ce-address 32.32.32.1
            exit
            sap 2/2/2:16 create
                ce-address 32.32.32.2
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-48>config>service#
```

Configuring PPP to Ethernet local Ipipe example:

Example:
```
config>service# ipipe 206 customer 1 create
config>service>ipipe$ sap 1/1/2:447 create
config>service>ipipe>sap$ description "eth_ppp_ipipe"
config>service>ipipe>sap$ ce-address 33.33.33.1
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# sap 2/2/2 create
config>service>ipipe>sap$ description "ppp_eth_ipipe"
config>service>ipipe>sap$ ce-address 33.33.33.2
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# no shutdown
config>service>ipipe# exit
config>service#
```

```
A:ALA-48>config>service# info
----------------------------------------------
...
        ipipe 206 customer 1 create
            sap 1/1/2:447 create
                description "eth_ppp_ipipe"
                ce-address 33.33.33.1
            exit
            sap 2/2/2 create
                description "ppp_eth_ipipe"
                ce-address 33.33.33.2
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-48>config>service#
```

Configuring Ipipe SDP Bindings

Creating a remote Ipipe configuration example:

Example:
```
config>service# sdp 16 mpls create
config>service>sdp$ far-end 4.4.4.4
config>service>sdp$ ldp
config>service>sdp$ path-mtu 1600
config>service>sdp$ no shutdown
config>service>sdp$ exit
config>service#

config>service# ipipe 207 customer 1 create
config>service>ipipe$ sap 1/1/2:449 create
config>service>ipipe>sap$ description "Remote_Ipipe"
config>service>ipipe>sap$ ce-address 34.34.34.1
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# spoke-sdp 16:516 create
config>service>ipipe>spoke-sdp$ ce-address 31.31.31.2
config>service>ipipe>spoke-sdp$ no shutdown
config>service>ipipe>spoke-sdp$ exit
config>service>ipipe#
```

```
config>service# info
----------------------------------------------
...
        sdp 16 mpls create
            far-end 4.4.4.4
            ldp
            path-mtu 1600
            keep-alive
                shutdown
            exit
            no shutdown
        exit
...
        ipipe 207 customer 1 create
            shutdown
            sap 1/1/2:449 create
                description "Remote_Ipipe"
                ce-address 34.34.34.1
            exit
            spoke-sdp 16:516 create
                ce-address 31.31.31.2
            exit
        exit
...

----------------------------------------------
A:ALA-48>config>service#
```

Using Spoke SDP Control Words

When control word is enabled, the Admin ControlWord is set to Preferred. It is necessary that both sides of the VLL have control word enabled or disabled for the pipe to be up.

The control word state will be set to True or False depending on what is configured on both sides of the VLL, either enabled (True) or disabled (False). When both sides are enabled (True) then the state is True. If one side is enabled (True) but the other side is disabled (False), then the state is False.

The command structure is the same for Ipipe and Apipe services. Apipe service data unit and Fpipes do not need to be specified.

Example:
```
config>service# epipe 2100 customer 1
        config>service>epipe$ description "Default epipe description
for service id 2100"
        config>service>epipe$ sap 1/2/7:4 create
        config>service>epipe>sap$ description "Default sap description
for service id 2100"
        config>service>epipe>sap$ exit
        config>service>epipe# spoke-sdp 1:2001 create
        config>service>epipe>spoke-sdp$ control-word
        config>service>epipe>spoke-sdp$ exit
        config>service>epipe# no shutdown
```

The following displays the configuration:

```
*A:ALA-Dut-B>config>service>epipe# info
----------------------------------------------
           description "Default epipe description for service id 2100"
           sap 1/2/7:4 create
               description "Default sap description for service id 2100"
           exit
           spoke-sdp 1:2001 create
               control-word
           exit
           no shutdown
----------------------------------------------
*A:ALA-Dut-B>config>service>epipe#
```

To disable the control word on spoke-sdp 1:2001:

Example:
```
config>service>epipe# spoke-sdp 1:2001 no control-word
    config>service>epipe>spoke-sdp$ exit

*A:ALA-Dut-B>config>service>epipe# info
----------------------------------------------
           description "Default epipe description for service id 2100"
           sap 1/2/7:4 create
               description "Default sap description for service id 2100"
           exit
           spoke-sdp 1:2001 create
```

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

See control-word on page 314 for show command output examples.

Pseudowire Configuration Notes

Note that the vc-switching parameter must be specified at the time of the VLL service creation.

Use the following CLI syntax to create pseudowire switching VLL services.

CLI Syntax: config>service# apipe *service-id* [customer *customer-id*] [vpn *vpn-id*] [vc-type {atm-vcc|atm-sdu|atm-vpc|atm-cell}] [vc-switching]
description *description-string*
spoke-sdp *sdp-id:vc-id*

CLI Syntax: config>service# epipe *service-id* [customer *customer-id*][vpn *vpn-id*][vc-switching]
description *description-string*
spoke-sdp *sdp-id:vc-id*

CLI Syntax: config>service# fpipe *service-id* [customer *customer-id*][vpn *vpn-id*] [vc-type {fr-dlci}] [vc-switching]
description *description-string*
spoke-sdp *sdp-id:vc-id*

CLI Syntax: config>service# ipipe *service-id* [customer *customer-id*][vpn *vpn-id*] [vc-switching]
description *description-string*
spoke-sdp *sdp-id:vc-id*

The following displays an example of the command usage to configure VLL pseudowire switching services:

Example:
```
config>service# apipe 1 customer 1 vpn 1 vc-switching create
       config>service>apipe$ description "Default apipe description for
service id 100"
       config>service>apipe# spoke-sdp 3:1 create
       config>service>apipe>spoke-sdp# exit
       config>service>apipe# spoke-sdp 6:200 create
       config>service>apipe>spoke-sdp# exit
       config>service>apipe# no shutdown
```

The following example displays configurations for each service:

```
*A:ALA-48>config>service# info
----------------------------------------------
...
       apipe 100 customer 1 vpn 1 vc-switching create
           description "Default apipe description for service id 100"
           spoke-sdp 3:1 create
           exit
           spoke-sdp 6:200 create
           exit
           no shutdown
       exit
...
```

```
        epipe 107 customer 1 vpn 107 vc-switching create
            description "Default epipe description for service id 107"
            spoke-sdp 3:8 create
            exit
            spoke-sdp 6:207 create
            exit
            no shutdown
        exit
...
        ipipe 108 customer 1 vpn 108 vc-switching create
            description "Default ipipe description for service id 108"
            spoke-sdp 3:9 create
            exit
            spoke-sdp 6:208 create
            exit
            no shutdown
        exit
...
        fpipe 109 customer 1 vpn 109 vc-switching creat9
            description "Default fpipe description for service id 108"
            spoke-sdp 3:5 create
            exit
            spoke-sdp 6:209 create
            exit
            no shutdown
        exit
...
-------------------------------------------
*A:ALA-48>config>service#
```

Configuring Two VLL Paths Terminating on T-PE2

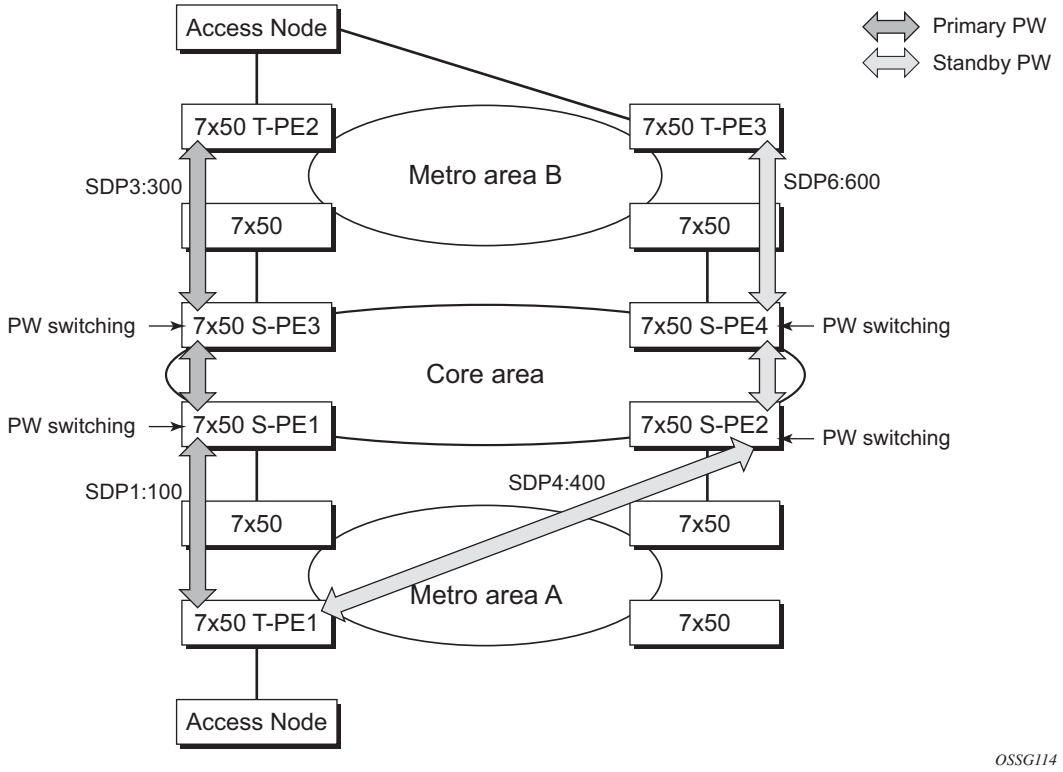

Figure 2: VLL Resilience with Pseudowire Redundancy and Switching

T-PE1

Example:

```
config>service# epipe 1 customer 1 vc-switching create
      config>service>epipe$ endpoint x create
      config>service>epipe>endpoint$ exit
      config>service>epipe# endpoint y create
      config>service>epipe>endpoint# revert-time 0
      config>service>epipe>endpoint# exit
      config>service>epipe# sap 1/1/1:100 endpoint x create
      config>service>epipe>sap$ exit
      config>service>epipe# spoke-sdp 1:100 endpoint y create
      config>service>epipe>spoke-sdp$ revert-timer 0
      config>service>epipe>spoke-sdp$ precedence primary
      config>service>epipe>spoke-sdp$ exit
      config>service>epipe# spoke-sdp 4:400 endpoint y create
      config>service>epipe>spoke-sdp$ precedence 0
      config>service>epipe>spoke-sdp$ no shutdown
      config>service>epipe>spoke-sdp$ exit
      config>service>epipe# no shutdown
```

The following example displays the configuration:

```
*A:ALA-T-PE1>config>service>epipe# info
----------------------------------------------
            endpoint "x" create
            exit
            endpoint "y" create
            exit
            sap 1/1/1:100 endpoint "x" create
            exit
            spoke-sdp 1:100 endpoint "y" create
                precedence primary
                revert-time 0
            exit
            spoke-sdp 4:400 endpoint "y" create
                precedence 0
            exit
            no shutdown
----------------------------------------------
*A:ALA-T-PE1>config>service>epipe#
```

T-PE2

Example:

```
config>service# epipe 1 customer 1 vc-switching create
config>service>epipe$ endpoint x create
config>service>epipe>endpoint$ exit
config>service>epipe# endpoint y create
config>service>epipe>endpoint# revert-time 0
config>service>epipe>endpoint# exit
config>service>epipe# sap 2/2/2:200 endpoint x create
config>service>epipe>sap$ exit
config>service>epipe# spoke-sdp 3:300 endpoint y primary create
config>service>epipe>spoke-sdp$ revert-timer 0
config>service>epipe>spoke-sdp$ exit
config>service>epipe# spoke-sdp 6:600 endpoint y create
config>service>epipe>spoke-sdp$ precedence 0
config>service>epipe>spoke-sdp$ no shutdown
config>service>epipe>spoke-sdp$ exit
config>service>epipe# no shutdown
```

The following example displays the configuration:

```
*A:ALA-T-PE2>config>service>epipe# info
----------------------------------------------
            endpoint "x" create
            exit
            endpoint "y" create
            exit
            sap 2/2/2:200 endpoint "x" create
            exit
            spoke-sdp 3:300 endpoint "y" create
                precedence primary
                revert-time 0
```

```
            exit
            spoke-sdp 6:600 endpoint "y" create
                precedence 0
            exit
            no shutdown
-----------------------------------------------
*A:ALA-T-PE2>config>service>epipe#
```

S-PE1: Note that specifying the **vc-switching** parameter enables a VC cross-connect so the service manager does not signal the VC label mapping immediately but will put this into passive mode.

Example:
```
config>service# epipe 1 customer 1 vc-switching create
        config>service>epipe# spoke-sdp 2:200 create
        config>service>epipe>spoke-sdp$ no shutdown
        config>service>epipe>spoke-sdp$ exit
        config>service>epipe# spoke-sdp 3:300 create
        config>service>epipe>spoke-sdp$ no shutdown
        config>service>epipe>spoke-sdp$ exit
        config>service>epipe# no shutdown
```

The following example displays the configuration:

```
*A:ALA-S-PE1>config>service>epipe# info
-----------------------------------------------
...
            spoke-sdp 2:200 create
            exit
            spoke-sdp 3:300 create
            exit
            no shutdown
-----------------------------------------------
*A:ALA-S-PE1>config>service>epipe#
```

S-PE2: Note that specifying the **vc-switching** parameter enables a VC cross-connect so the service manager does not signal the VC label mapping immediately but will put this into passive mode.

Example:
```
config>service# epipe 1 customer 1 vc-switching create
        config>service>epipe# spoke-sdp 2:200 create
        config>service>epipe>spoke-sdp$ no shutdown
        config>service>epipe>spoke-sdp$ exit
        config>service>epipe# spoke-sdp 3:300 create
        config>service>epipe>spoke-sdp$ no shutdown
        config>service>epipe>spoke-sdp$ exit
        config>service>epipe# no shutdown
```

The following example displays the configuration:

```
*A:ALA-S-PE2>config>service>epipe# info
----------------------------------------------
...
            spoke-sdp 2:200 create
            exit
            spoke-sdp 3:300 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-S-PE2>config>service>epipe#
```

Configuring VLL Resilience

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

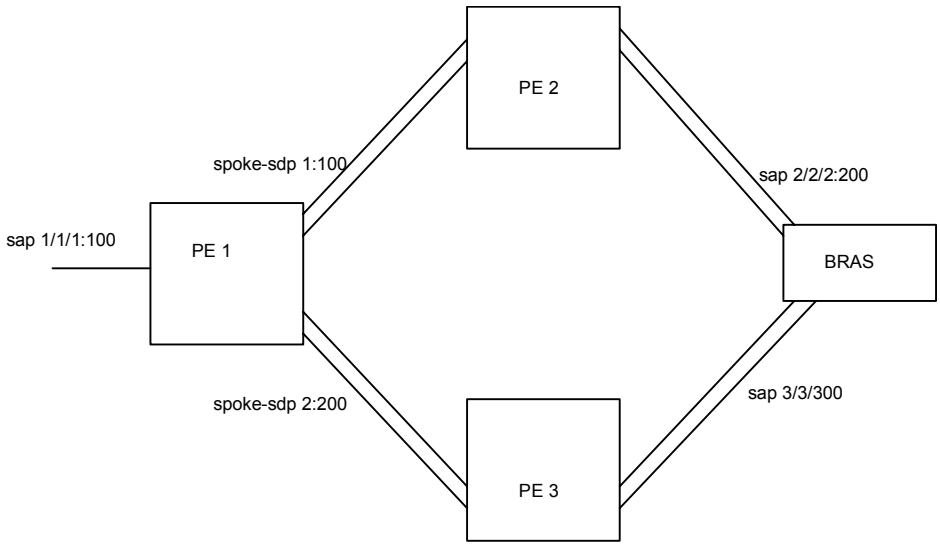

Figure 3: VLL Resilience

For PE1:

Example:
```
config>service# epipe 12 customer 1 vc-switching create
        config>service>epipe$ endpoint x create
        config>service>epipe>endpoint$ exit
        config>service>epipe# endpoint y create
        config>service>epipe>endpoint# revert-time 0
        config>service>epipe>endpoint# exit
        config>service>epipe# sap 1/1/1:100 endpoint x create
        config>service>epipe>sap$ exit
        config>service>epipe# spoke-sdp 1:100 endpoint y create
        config>service>epipe>spoke-sdp$ precedence primary
        config>service>epipe>spoke-sdp$ exit
        config>service>epipe# spoke-sdp 2:200 endpoint y create
        config>service>epipe>spoke-sdp$ precedence 1
        config>service>epipe>spoke-sdp$ no shutdown
        config>service>epipe>spoke-sdp$ exit
        config>service>epipe# no shutdown
```

The following example displays the configuration:

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

See details on the rules for the determination of the active transmit endpoint object based on endpoint status transitions in VLL Endpoint Active Transmit Object Selection Rules on page 139.

Configuring VLL Resilience for a Switched Pseudowire Path

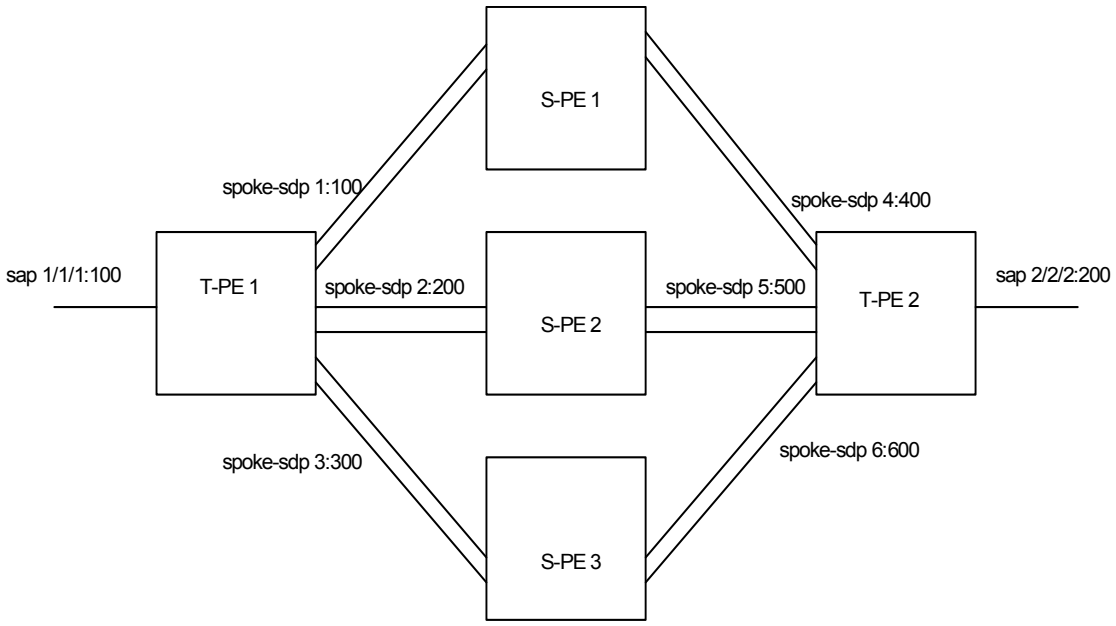

Figure 4: VLL Resilience with PW Switching

T-PE1

Example:
```
config>service# epipe 1 customer 1 vc-switching create
        config>service>epipe$ endpoint x create
        config>service>epipe>endpoint$ exit
        config>service>epipe# endpoint y create
        config>service>epipe>endpoint# revert-time 0
        config>service>epipe>endpoint# exit
        config>service>epipe# sap 1/1/1:100 endpoint x create
        config>service>epipe>sap$ exit
        config>service>epipe# spoke-sdp 1:100 endpoint y  create
        config>service>epipe>spoke-sdp$ precedence primary
        config>service>epipe>spoke-sdp$ exit
        config>service>epipe# spoke-sdp 2:200 endpoint y  create
        config>service>epipe>spoke-sdp$ precedence 1
        config>service>epipe>spoke-sdp$ no shutdown
        config>service>epipe>spoke-sdp$ exit
        config>service>epipe# spoke-sdp 3:300 endpoint y create
        config>service>epipe>spoke-sdp$ precedence 1
```

```
        config>service>epipe>spoke-sdp$ exit
        config>service>epipe# no shutdown
```

The following example displays the configuration:

```
*A:ALA-48>config>service>epipe# info
----------------------------------------------
            endpoint "x" create
            exit
            endpoint "y" create
            exit
            sap 1/1/1:100 endpoint "x" create
            exit
            spoke-sdp 1:100 endpoint "y" create
                precedence primary
            exit
            spoke-sdp 2:200 endpoint "y" create
                precedence 1
            exit
            spoke-sdp 3:300 endpoint "y" create
                precedence 1
            exit
            no shutdown
----------------------------------------------
*A:ALA-48>config>service>epipe#
```

T-PE2

Example:

```
config>service# epipe 1 customer 1 vc-switching create
        config>service>epipe$ endpoint x create
        config>service>epipe>endpoint$ exit
        config>service>epipe# endpoint y create
        config>service>epipe>endpoint# revert-time 100
        config>service>epipe>endpoint# exit
        config>service>epipe# sap 2/2/2:200 endpoint x create
        config>service>epipe>sap$ exit
        config>service>epipe# spoke-sdp 4:400 endpoint y  create
        config>service>epipe>spoke-sdp$ precedence primary
        config>service>epipe>spoke-sdp$ exit
        config>service>epipe# spoke-sdp 5:500 endpoint y  create
        config>service>epipe>spoke-sdp$ precedence 1
        config>service>epipe>spoke-sdp$ no shutdown
        config>service>epipe>spoke-sdp$ exit
        config>service>epipe# spoke-sdp 6:600 endpoint y  create
        config>service>epipe>spoke-sdp$ precedence 1
        config>service>epipe>spoke-sdp$ exit
        config>service>epipe# no shutdown
```

The following example displays the configuration:

```
*A:ALA-49>config>service>epipe# info
----------------------------------------------
            endpoint "x" create
            exit
            endpoint "y" create
                revert-time 100
            exit
            sap 2/2/2:200 endpoint "x" create
            exit
            spoke-sdp 4:400 endpoint "y" create
                precedence primary
            exit
            spoke-sdp 5:500 endpoint "y" create
                precedence 1
            exit
            spoke-sdp 6:600 endpoint "y" create
                precedence 1
            exit
            no shutdown
----------------------------------------------
*A:ALA-49>config>service>epipe#
```

S-PE1

Example:
```
config>service# epipe 1 customer 1 vc-switching create
config>service>epipe# spoke-sdp 1:100 create
config>service>epipe>spoke-sdp$ no shutdown
config>service>epipe>spoke-sdp$ exit
config>service>epipe# spoke-sdp 4:400 create
config>service>epipe>spoke-sdp$ no shutdown
config>service>epipe>spoke-sdp$ exit
config>service>epipe# no shutdown
```

The following example displays the configuration:

```
*A:ALA-50>config>service>epipe# info
----------------------------------------------
...
            spoke-sdp 1:100 create
            exit
            spoke-sdp 4:400 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-49>config>service>epipe#
```

See details on the rules for the determination of the active transmit endpoint object based on endpoint status transitions in VLL Endpoint Active Transmit Object Selection Rules on page 139.

Service Management Tasks

This section discusses the following Epipe service management tasks:

This section discusses the following Apipe service management tasks:

This section discusses the following Fpipe service management tasks:

This section discusses the following Ipipe service management tasks:

Modifying Epipe Service Parameters

The following displays an example of adding an accounting policy to an existing SAP:

Example:
```
config>service# epipe 2
       config>service>epipe# sap 2/1/3:21
       config>service>epipe>sap# accounting-policy 14
       config>service>epipe>sap# exit
```

The following output displays the SAP configuration:

```
ALA-1>config>service# info
----------------------------------------------
    epipe 2 customer 6 vpn 2 create
           description "Distributed Epipe service to east coast"
           sap 2/1/3:21 create
             accounting-policy 14
           exit
           spoke-sdp 2:6000 create
           exit
           no shutdown
       exit
----------------------------------------------
ALA-1>config>service#
```

Disabling an Epipe Service

You can shut down an Epipe service without deleting the service parameters.

CLI Syntax: `config>service> epipe` *service-id*
`shutdown`

Example:
```
config>service# epipe 2
       config>service>epipe# shutdown
       config>service>epipe# exit
```

Re-enabling an Epipe Service

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

CLI Syntax:
```
config>service> epipe service-id
            no shutdown
```

Example:
```
config>service# epipe 2
       config>service>epipe# no shutdown
       config>service>epipe# exit
```

Deleting an Epipe Service

Perform the following steps prior to deleting an Epipe service:

1. Shut down the SAP and SDP.
2. Delete the SAP and SDP.
3. Shut down the service.

Use the following CLI syntax to delete an Epipe service:

CLI Syntax:
```
config>service
            [no] epipe service-id
               shutdown
               [no] sap sap-id
                  shutdown
               [no] spoke-sdp sdp-id:vc-id
                  shutdown
```

Example:
```
config>service# epipe 2
       config>service>epipe# sap 2/1/3:21
       config>service>epipe>sap# shutdown
       config>service>epipe>sap# exit
       config>service>epipe# no sap 2/1/3:21
       config>service>epipe# spoke-sdp 2:6000
       config>service>epipe>spoke-sdp# shutdown
       config>service>epipe>spoke-sdp# exit
       config>service>epipe# no spoke-sdp 2:6000
       config>service>epipe# epipe 2
       config>service>epipe# shutdown
       config>service>epipe# exit
       config>service# no epipe 2
```

Modifying Apipe Service Parameters

The following example displays command usage to modify Apipe parameters:

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# sap 1/1/1:0/32 create
A:ALA-41>config>service>apipe>sap# accounting-policy 2
A:ALA-41>config>service>apipe>sap# exit
A:ALA-41>config>service>apipe# spoke-sdp 1:4
A:ALA-41>config>service>apipe>spoke-sdp# egress
A:ALA-41>config>service>apipe>spoke-sdp>egress# vc-label 16
A:ALA-41>config>service>apipe>spoke-sdp>egress# exit
A:ALA-41>config>service>apipe>spoke-sdp# exit
A:ALA-41>config>service>apipe#
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# apipe 5
A:ALA-42>config>service>apipe# sap 2/2/2:0/32 create
A:ALA-42>config>service>apipe>sap# accounting-policy 2
A:ALA-42>config>service>apipe>sap# exit
A:ALA-42>config>service>apipe# spoke-sdp 1:4
A:ALA-42>config>service>apipe>spoke-sdp# egress
A:ALA-42>config>service>apipe>spoke-sdp>egress# vc-label 16
A:ALA-42>config>service>apipe>spoke-sdp>egress# exit
A:ALA-42>config>service>apipe>spoke-sdp# exit
A:ALA-42>config>service>apipe#
```

```
PE Router 1 (ALA-41):
A:ALA-41>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 1/1/1:0/32 create
               accounting-policy 2
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:4 create
               egress
                   vc-label 16
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-41>config>service#
```

```
PE Router 2 (ALA-42):
A:ALA-42>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 2/2/2:0/32 create
               accounting-policy 2
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:4 create
               egress
                   vc-label 16
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-42>config>service#
```

Disabling an Apipe Service

An Apipe service can be shut down without deleting any service parameters.

CLI Syntax:
```
config>service#
        apipe service-id
           shutdown
```

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# shutdown
A:ALA-41>config>service>apipe# exit
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# apipe 5
A:ALA-42>config>service>apipe# shutdown
A:ALA-42>config>service>apipe# exit
```

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
----------------------------------
...
        apipe 5 customer 1 create
            shutdown
            description "apipe test"
            service-mtu 1400
            sap 1/1/1:0/32 create
               accounting-policy 2
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:4 create
               egress
                   vc-label 16
            exit
            no shutdown
        exit
...
----------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
----------------------------------
...
        apipe 5 customer 1 create
            shutdown
            description "apipe test"
            service-mtu 1400
            sap 2/2/2:0/32 create
               accounting-policy 2
               ingress
```

```
                qos 102
             exit
             egress
                qos 103
             exit
          exit
          spoke-sdp 1:4 create
             egress
                vc-label 16
             exit
       exit
...
------------------------------------
A:ALA-42>config>service#
```

Re-enabling an Apipe Service

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

CLI Syntax: `config>service#`
`apipe` *`service-id`*
`no shutdown`

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# no shutdown
A:ALA-41>config>service>apipe# exit
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# apipe 5
A:ALA-42>config>service>apipe# no shutdown
A:ALA-42>config>service>apipe# exit
```

Deleting an Apipe Service

An Apipe service cannot be deleted until the SAP is shut down. 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 Apipe services:

CLI Syntax:
```
config>service#
      no apipe service-id
         shutdown
         no sap sap-id
            shutdown
         no spoke-sdp [sdp-id:vc-id]
            shutdown
```

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# sap 1/1/1:0/32
A:ALA-41>config>service>apipe>sap# shutdown
A:ALA-41>config>service>apipe>sap# exit
A:ALA-41>config>service>apipe# no sap 1/1/1:0/32
A:ALA-41>config>service>apipe# spoke-sdp 1:4
A:ALA-41>config>service>apipe>spoke-sdp# shutdown
A:ALA-41>config>service>apipe>spoke-sdp# exit
A:ALA-41>config>service>apipe# no spoke-sdp 1:4
A:ALA-41>config>service>apipe# shutdown
A:ALA-41>config>service>apipe# exit
A:ALA-41>config>service# no apipe 5
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-41>config>service# apipe 5
A:ALA-41>config>service>apipe# sap 2/2/2:0/32
A:ALA-41>config>service>apipe>sap# shutdown
A:ALA-41>config>service>apipe>sap# exit
A:ALA-41>config>service>apipe# no sap 2/2/2:0/32
A:ALA-41>config>service>apipe# spoke-sdp 1:4
A:ALA-41>config>service>apipe>spoke-sdp# shutdown
A:ALA-41>config>service>apipe>spoke-sdp# exit
A:ALA-41>config>service>apipe# no spoke-sdp 1:4
A:ALA-41>config>service>apipe# shutdown
A:ALA-41>config>service>apipe# exit
A:ALA-41>config>service# no apipe 5
```

Modifying Fpipe Service Parameters

The following example displays command usage to modify Fpipe parameters:

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# sap 1/2/1:16 create
A:ALA-41>config>service>fpipe>sap# accounting-policy 2
A:ALA-41>config>service>fpipe>sap# exit
A:ALA-41>config>service>fpipe# spoke-sdp 1:4
A:ALA-41>config>service>fpipe>spoke-sdp# ingress
A:ALA-41>config>service>fpipe>spoke-sdp>filter ip 10
A:ALA-41>config>service>fpipe>spoke-sdp# exit
A:ALA-41>config>service>fpipe#
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# fpipe 1
A:ALA-42>config>service>fpipe# sap 2/1/1.1:16 create
A:ALA-42>config>service>fpipe>sap# accounting-policy 2
A:ALA-42>config>service>fpipe>sap# exit
A:ALA-42>config>service>fpipe# spoke-sdp 1:1
A:ALA-42>config>service>fpipe>spoke-sdp# egress
A:ALA-42>config>service>fpipe>spoke-sdp>egress# filter ip 10
A:ALA-42>config>service>fpipe>spoke-sdp>egress# exit
A:ALA-42>config>service>fpipe>spoke-sdp# exit
A:ALA-42>config>service>fpipe#
```

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
------------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 1/2/1:16 create
               accounting-policy 2
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
               ingress
                   filter ip 10
            exit
            no shutdown
        exit
...
------------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
----------------------------------
...
        fpipe 1 customer 1 create
            description "fpipe test"
            service-mtu 1400
            sap 2/1/1.1:16 create
               accounting-policy 2
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
               egress
                   filter ip 10
            exit
            no shutdown
        exit
...
----------------------------------
A:ALA-42>config>service#
```

Disabling an Fpipe Service

An Fpipe service can be shut down without deleting any service parameters.

CLI Syntax:
```
config>service#
    fpipe service-id
        shutdown
```

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# shutdown
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# fpipe 1
A:ALA-42>config>service>fpipe# shutdown
```

PE Router 1 (ALA-41):

```
A:ALA-41>config>service# info
------------------------------------
...
        fpipe 1 customer 1 create
            shutdown
            description "fpipe test"
            service-mtu 1400
            sap 1/2/1:16 create
               accounting-policy 2
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
               ingress
                   filter ip 10
            exit
        exit
...
------------------------------------
A:ALA-41>config>service#
```

PE Router 2 (ALA-42):

```
A:ALA-42>config>service# info
------------------------------------
...
        fpipe 1 customer 1 create
```

```
            shutdown
            description "fpipe test"
            service-mtu 1400
            sap 2/1/1.1:16 create
               accounting-policy 2
               ingress
                   qos 101
               exit
               egress
                   qos 1020
               exit
            exit
            spoke-sdp 1:1 create
               egress
                   filter ip 10
            exit
        exit
...
-----------------------------------
A:ALA-42>config>service#
```

Re-enabling an Fpipe Service

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

CLI Syntax: `config>service#`
`fpipe` *`service-id`*
`no shutdown`

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# no shutdown
A:ALA-41>config>service>fpipe# exit
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# fpipe 1
A:ALA-42>config>service>fpipe# no shutdown
A:ALA-42>config>service>fpipe# exit
```

Deleting an Fpipe Service

An Fpipe service cannot be deleted until the SAP is shut down. 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 Fpipe service:

CLI Syntax:
```
config>service#
      no fpipe service-id
         shutdown
         no sap sap-id
            shutdown
         no spoke-sdp [sdp-id:vc-id]
            shutdown
```

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# sap 1/1/1:0/32
A:ALA-41>config>service>fpipe>sap# shutdown
A:ALA-41>config>service>fpipe>sap# exit
A:ALA-41>config>service>fpipe# no sap 1/1/1:0/32
A:ALA-41>config>service>fpipe# spoke-sdp 1:1
A:ALA-41>config>service>fpipe>spoke-sdp# shutdown
A:ALA-41>config>service>fpipe>spoke-sdp# exit
A:ALA-41>config>service>fpipe# no spoke-sdp 1:1
A:ALA-41>config>service>fpipe# shutdown
A:ALA-41>config>service>fpipe# exit
A:ALA-41>config>service# no fpipe 1
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-41>config>service# fpipe 1
A:ALA-41>config>service>fpipe# sap 2/1/1.1:16
A:ALA-41>config>service>fpipe>sap# shutdown
A:ALA-41>config>service>fpipe>sap# exit
A:ALA-41>config>service>fpipe# no sap 2/1/1.1:16
A:ALA-41>config>service>fpipe# spoke-sdp 1:1
A:ALA-41>config>service>fpipe>spoke-sdp# shutdown
A:ALA-41>config>service>fpipe>spoke-sdp# exit
A:ALA-41>config>service>fpipe# no spoke-sdp 1:1
A:ALA-41>config>service>fpipe# shutdown
A:ALA-41>config>service>fpipe# exit
A:ALA-41>config>service# no fpipe 1
```

Modifying Ipipe Service Parameters

The following example displays command usage to modify Ipipe parameters:

Example:
```
config>service# ipipe 202
config>service>ipipe# sap 1/1/2:444
config>service>ipipe>sap# shutdown
config>service>ipipe>sap# exit
config>service>ipipe# no sap 1/1/2:444
config>service>ipipe# sap 1/1/2:555 create
config>service>ipipe>sap$ description "eth_ipipe"
config>service>ipipe>sap$ ce-address 31.31.31.1
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# info
```

```
A:ALA-48>config>service# info
----------------------------------------------
...
        ipipe 202 customer 1 create
            sap 1/1/2:445 create
                description "eth_ipipe"
                ce-address 31.31.31.2
            exit
            sap 1/1/2:555 create
                description "eth_ipipe"
                ce-address 31.31.31.1
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-48>config>service#
```

Disabling an Ipipe Service

An Ipipe service can be shut down without deleting any service parameters.

CLI Syntax:
```
config>service#
    ipipe service-id
        shutdown
```

Example:
```
A:ALA-41>config>service# ipipe 202
A:ALA-41>config>service>ipipe# shutdown
```

```
A:ALA-48>config>service# info
----------------------------------------------
...
       ipipe 202 customer 1 create
           shutdown
           sap 1/1/2:445 create
               description "eth_ipipe"
               ce-address 31.31.31.2
           exit
           sap 1/1/2:555 create
               description "eth_ipipe"
               ce-address 31.31.31.1
           exit
       exit
...
----------------------------------------------
A:ALA-48>config>service#
```

Re-enabling an Ipipe Service

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

CLI Syntax:
```
config>service#
    ipipe service-id
        no shutdown
```

Example:
```
A:ALA-41>config>service# ipipe 202
A:ALA-41>config>service>ipipe# no shutdown
```

Deleting an Ipipe Service

An Ipipe service cannot be deleted until the SAP is shut down. 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 an Ipipe service:

CLI Syntax:
```
config>service#
no ipipe service-id
   shutdown
   no sap sap-id
      shutdown
   no spoke-sdp [sdp-id:vc-id]
      shutdown
```

Example:
```
config>service# ipipe 207
config>service>ipipe# sap 1/1/2:449
config>service>ipipe>sap# shutdown
config>service>ipipe>sap# exit
config>service>ipipe# no sap 1/1/2:449
config>service>ipipe# spoke-sdp 16:516
config>service>ipipe>spoke-sdp# shutdown
config>service>ipipe>spoke-sdp# exit
config>service>ipipe# no  spoke-sdp 16:516
config>service>ipipe# exit
config>service# no ipipe 207
config>service#
```

VLL Services Command Reference

Command Hierarchies

Apipe Service Configuration Commands

```
config
    — service
        — apipe service-id [customer customer-id] [vpn vpn-id] [vc-type {atm-vcc | atm-sdu | atm-vpc
            | atm-cell}] [vc-switching]
        — no apipe service-id
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — active-hold-delay active-endpoint-delay
                — no active-hold-delay
                — description description-string
                — no description
                — revert-time  revert-time
                — no revert-time
            — interworking {frf-5}
            — no interworking
            — sap sap-id [no-endpoint]
            — sap sap-id [endpoint endpoint-name]
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy
                — atm
                    — egress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — ingress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — oam
                        — [no] alarm-cells
                        — [no] terminate
                — [no] collect-stats
                — description description-string
                — no description
                — egress
```

```
                    — agg-rate-limit agg-rate
                    — no agg-rate-limit
                    — [no] qinq-mark-top-only
                    — qos policy-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
                            — rate  pir-rate [cir cir-rate]
                            — no rate
                    — scheduler-policy scheduler-policy-name
                    — no scheduler-policy
                — ingress
                    — qos policy-id [shared-queuing]
                    — no qos
                    — [no] queue-override
                        — [no] queue queue-id
                            — adaptation-rule [pir adaptation-rule] [cir
                              adaptation-rule]
                            — no adaptation-rule
                            — cbs size-in-kbytes
                            — no cbs
                            — 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
                            — rate  pir-rate [cir cir-rate]
                            — no rate
                    — scheduler-policy scheduler-policy-name
                    — no scheduler-policy
                — multi-service-site customer-site-name
                — no multi-service-site
                — [no] shutdown
                — tod-suite tod-suite-name
                — no tod-suite
            — service-mtu octets
            — no service-mtu
            — [no] shutdown
            — spoke-sdp  [sdp-id[:vc-id] [no-endpoint]
```

— **spoke-sdp** [*sdp-id*[:*vc-id*] **endpoint** *endpoint-name* [**icb**]
— **no spoke-sdp** [*sdp-id*[:*vc-id*]
    — **cell-concatenation**
        — [**no**] **aal5-frame-aware**
        — [**no**] **clp-change**
        — **max-cells** *cell-count*
        — **no max-cellss** [*cell-count*>]
        — **max-delay** *delay-time*
        — **no max-delay** [*delay-time*]
    — [**no**] **control-word**
    — **egress**
        — **vc-label** *ingress-vc-label*
        — **no vc-label** [*ingress-vc-label*]
    — **ingress**
        — **vc-label** *ingress-vc-label*
        — **no vc-label** [*ingress-vc-label*]
    — **precedence** [*precedence-value*| **primary**]
    — **no precedence**
    — [**no**] **shutdown**

Epipe Service Configuration Commands

```
config
    — service
        — [no] epipe service-id [customer customer-id] [vpn vpn-id] [vc-switching]
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — active-hold-delay active-endpoint-delay
                — no active-hold-delay
                — description description-string
                — no description
                — revert-time [revert-time | infinite]
                — no revert-time
            — sap sap-id [no-endpoint]
            — sap sap-id [endpoint endpoint-name]
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy
                — atm
                    — egress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — encapsulation atm-encap-type
                    — ingress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — oam
                        — [no] alarm-cells
                — [no] collect-stats
                — description description-string
                — no description
                — egress
                    — agg-rate-limit agg-rate
                    — no agg-rate-limit
                    — filter [ip ip-filter-id]
                    — filter [ipv6 ipv6-filter-id]
                    — filter [mac mac-filter-id]
                    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
                    — [no] qinq-mark-top-only
                    — qos policy-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
                                            — rate pir-rate [cir cir-rate]
                                            — no rate
                                    — scheduler-policy scheduler-policy-name
                                    — no scheduler-policy
                            — ingress
                                    — filter [ip ip-filter-id]
                                    — filter [ipv6 ipv6-filter-id]
                                    — filter [mac mac-filter-id]
                                    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
                                    — match-qinq-dot1p {top | bottom}
                                    — no match-qinq-dot1p
                                    — qos policy-id [shared-queuing]
                                    — no qos
                                    — [no] queue-override
                                        — [no] queue queue-id
                                            — adaptation-rule [pir adaptation-rule] [cir
                                              adaptation-rule]
                                            — no adaptation-rule
                                            — cbs size-in-kbytes
                                            — no cbs
                                            — high-prio-only percent
                                            — no high-prio-only
                                            — mbs size-in-kbytes
                                            — no mbs
                                            — rate pir-rate [cir cir-rate]
                                            — no rate
                                    — [no] scheduler-override
                                        — [no] scheduler scheduler-name
                                            — rate pir-rate [cir cir-rate]
                                            — no rate
                                    — scheduler-policy scheduler-policy-name
                                    — no scheduler-policy
                            — multi-service-site customer-site-name
                            — no multi-service-site
                            — tod-suite tod-suite-name
                            — no tod-suite
                            — [no] shutdown
                    — service-mtu octets
                    — no service-mtu
                    — [no] shutdown
                    — spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [no-endpoint]
                    — spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] endpoint [icb]
                    — no spoke-sdp sdp-id[:vc-id]
                            — accounting-policy acct-policy-id
                            — no accounting-policy
                            — [no] collect-stats
                            — [no] control-word
                            — [no] egress
                                    — filter [ip ip-filter-id]
                                    — filter [ipv6 ipv6-filter-id]
                                    — filter [mac mac-filter-id]
                                    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
```

— [**no**] **vc-label** *egress-vc-label*
— [**no**] **ingress**
— **filter** [**ip** *ip-filter-id*]
— **filter** [**ipv6** *ipv6-filter-id*]
— **filter** [**mac** *mac-filter-id*]
— **no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*] [**mac** *mac-filter-id*]
— [**no**] **vc-label** *ingress-vc-label*
— **precedence** [*precedence-value*| **primary**]
— **no precedence**
— [**no**] **shutdown**
— **vlan-vc-tag** *0..4094*
— **no vlan-vc-tag** [*0..4094*]

Fpipe Service Configuration Commands

```
config
    — service
        — fpipe service-id [customer customer-id] [vpn vpn-id] [vc-type {fr-dlci}] [vc-switching]
        — no fpipe service-id
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — active-hold-delay active-endpoint-delay
                — no active-hold-delay
                — description description-string
                — no description
                — revert-time revert-time
                — no revert-time
            — sap sap-id [no-endpoint]
            — sap sap-id endpoint endpoint-name
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy
                — [no] collect-stats
                — description description-string
                — no description
                — egress
                    — agg-rate-limit agg-rate
                    — no agg-rate-limit
                    — filter [ip ip-filter-id]
                    — filter [ipv6 ipv6-filter-id]
                    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                    — [no] qinq-mark-top-only
                    — qos policy-id
                    — no qos
                    — [no] queue-override
                        — [no] queue queue-id
                            — adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
                            — no adaptation-rule
                            — avg-frame-overhead percent
                            — 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
                            — rate pir-rate [cir cir-rate]
                            — no rate
                    — scheduler-policy scheduler-policy-name
                    — no scheduler-policy
                — ingress
                    — filter [ip ip-filter-id]
```

```
                    — filter [ipv6 ipv6-filter-id]
                    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                    — qos policy-id [shared-queuing]
                    — no qos
                    — [no] queue-override
                        — [no] queue queue-id
                            — adaptation-rule [pir adaptation-rule] [cir
                              adaptation-rule]
                            — no adaptation-rule
                            — avg-frame-overhead percent
                            — 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
                            — rate  pir-rate [cir cir-rate]
                            — no rate
                    — scheduler-policy scheduler-policy-name
                    — no scheduler-policy
                    — scheduler-policy scheduler-policy-name
                    — no scheduler-policy
                — multi-service-site customer-site-name
                — no multi-service-site
                — [no] shutdown
                — tod-suite tod-suite-name
                — no tod-suite
            — service-mtu octets
            — no service-mtu
            — [no] shutdown
            — spoke-sdp  sdp-id[:vc-id] [no-endpoint]
            — spoke-sdp  sdp-id[:vc-id] endpoint endpoint-name [icb]
            — no spoke-sdp  sdp-id[:vc-id]
                — egress
                    — filter [ip ip-filter-id]
                    — filter [ipv6 ipv6-filter-id]
                    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                    — vc-label ingress-vc-label
                    — no vc-label [ingress-vc-label]
                — ingress
                    — filter [ip ip-filter-id]
                    — filter [ipv6 ipv6-filter-id]
                    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                    — vc-label ingress-vc-label
                    — no vc-label [ingress-vc-label]
                — precedence  [precedence-value| primary]
                — no precedence
                — [no] shutdown
```

Ipipe Service Configuration Commands

```
config
    — service
        — ipipe service-id [customer customer-id] [vpn vpn-id] [vc-switching]
        — no ipipe service-id
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — active-hold-delay active-endpoint-delay
                — no active-hold-delay
                — description description-string
                — no description
                — revert-time  revert-time
                — no revert-time
            — sap sap-id [no-endpoint]
            — sap sap-id endpoint endpoint-name
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy
                — atm
                    — egress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — encapsulation atm-encap-type
                    — ingress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — oam
                        — [no] alarm-cells
                — ce-address ip-address
                — no ce-address
                — collect-stats
                — no collect-stats
                — description description-string
                — no description
                — egress
                    — agg-rate-limit agg-rate
                    — no agg-rate-limit
                    — filter [ip ip-filter-id]
                    — no filter
                    — [no] qinq-mark-top-only
                    — qos policy-id
                    — no qos
                    — [no] queue-override
                        — [no] queue queue-id
                            — adaptation-rule [pir adaptation-rule] [cir
                              adaptation-rule]
                            — no adaptation-rule
                            — avg-frame-overhead percent
                            — 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
                                            — rate  pir-rate [cir cir-rate]
                                            — no rate
                                — scheduler-policy scheduler-policy-name
                                — no scheduler-policy
                        — ingress
                                — filter [ip ip-filter-id]
                                — no filter
                                — match-qinq-dot1p {top | bottom}
                                —  no match-qinq-dot1p
                                — qos policy-id [shared-queuing]
                                — no qos
                                — [no] queue-override
                                        — [no] queue queue-id
                                            — adaptation-rule [pir adaptation-rule] [cir
                                              adaptation-rule]
                                            — no adaptation-rule
                                            — cbs size-in-kbytes
                                            — no cbs
                                            — 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
                                            — rate  pir-rate [cir cir-rate]
                                            — no rate
                                — scheduler-policy scheduler-policy-name
                                — no scheduler-policy
                        — mac [ieee-address]
                        — no mac
                        — mac-refresh [refresh interval]
                        — no mac-refresh
                        — multi-service-site customer-site-name
                        — no multi-service-site
                        — [no] shutdown
                        — tod-suite tod-suite-name
                        — no tod-suite
        — service-mtu octets
        — no service-mtu
        — [no] shutdown
        — spoke-sdp [sdp-id[:vc-id] [no-endpoint]
        — spoke-sdp [sdp-id[:vc-id] endpoint endpoint-name [icb]
        — no spoke-sdp sap-id
                        — ce-address ip-address
                        — no ce-address
                        — [no] control-word
                        — egress
                                — filter {ip ip-filter-id }
```

```
            — no filter
            — [no] vc-label  vc-label
— ingress
            — filter {ip ip-filter-id}
            — no filter
            — vc-label  ingress-vc-label
            — no vc-label [ingress-vc-label]
— precedence  [precedence-value| primary]
— no precedence
— [no] shutdown
```

Show Commands

```
show
    — service
        — egress-label start-label [end-label]
        — ingress-label start-label [end-label]
        — 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
        — sdp [sdp-id | far-end ip-addr] [detail | keep-alive-history]
        — sdp-using [sdp-id[:vc-id] | far-end ip-address]
        — service-using [epipe] [ies] [vpls] [vprn] [mirror] [apipe] [fpipe] [ipipe] [sdp sdp-id]
          [customer customer-id]
        — id service-id
            — all
            — authentication
                — statistics [policy name] [sap sap-id]
            — base
            — endpoint [endpoint-name]
            — labels
            — lease-state [[sap sap-id]  [sdp [sdp-id[:vc-id]]] | [interface interface-name] | [ip-
              address ip-address[/mask]] | [mac ieee-address] | [wholesaler service-id] ] [detail]
            — retailers
            — sap [sap-id] [detail]]
            — sdp [sdp-id | far-end ip-address] [detail]
            — wholesalers
```

Clear Commands

```
clear
    — service
        — id service-id
            — fdb {all | mac ieee-address | sap sap-id | mesh-sdp sdp-id[:vc-id] | spoke-sdp sdp-
              id:vc-id}
            — spoke-sdp sdp-id:vc-id ingress-vc-label
        — statistics
            — id service-id
                — counters
                — spoke-sdp sdp-id:vc-id {all | counters | stp}
            — sap sap-id {all | counters | stp}
            — sdp sdp-id keep-alive
```

VLL Service Configuration Commands

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>service>epipe
config>service>epipe>sap
config>service>epipe>spoke-sdp
config>service>apipe
config>service>apipe>sap
config>service>apipe>spoke-sdp
config>service>fpipe
config>service>fpipe>sap
config>service>fpipe>spoke-sdp
config>service>ipipe
config>service>ipipe>sap
config>service>ipipe>spoke-sdp

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

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

Service Operational State — A service is regarded as operational providing that at least one SAP and one SDP are operational or if two SAP's are operational.

SDP (global) — When an SDP is shutdown at the global service level, all bindings to that SDP are put into the out-of-service state and the SDP itself is put into the administratively and operationally down states. Packets that would normally be transmitted using this SDP binding will be discarded and counted as dropped packets.

SDP (service level) — Shutting down an SDP within a service only affects traffic on that service from entering or being received from the SDP. The SDP itself may still be operationally up for other services.

description

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

Context config>service>epipe
config>service>epipe>sap
config>service>epipe>endpoint
config>service>apipe
config>service>apipe>sap
config>service>apipe>endpoint
config>service>fpipe
config>service>fpipe>sap
config>service>fpipe>endpoint
config>service>ipipe
config>service>ipipe>sap
config>service>ipipe>endpoint

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

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

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

Default No description associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters 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.

VLL Global Commands

epipe

Syntax **epipe** *service-id* **customer** *customer-id* [**vpn** *vpn-id*] [**vc-switching**]
epipe *service-id*
no epipe *service-id*

Context config>service

Description This command configures an epipe service instance. This command is used to configure a point-to-point epipe service. An epipe connects two endpoints defined as Service Access Points (SAPs). Both SAPs may be defined in one 7750 SR or they may be defined in separate 7750 SR devices connected over the service provider network. When the endpoint SAPs are separated by the service provider network, the far end SAP is generalized into a Service Distribution Point (SDP). This SDP describes a destination 7750 SR and the encapsulation method used to reach it.

No MAC learning or filtering is provided on an epipe.

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

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

By default, no epipe services exist until they are explicitly created with this command.

The **no** form of this command deletes the epipe service instance with the specified *service-id*. The service cannot be deleted until the service has been shutdown.

service-id — The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR on which this service is defined.

Values 1 — 2147483647

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 ID. If this parameter is not specified, the VPN ID uses the same service ID number.

Values 1 — 2147483647

Default **null** (0)

vc-switching — Specifies if the pseudowire switching signalling is used for the spoke SDPs configured in this service.

apipe

Syntax **apipe** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*] [**vc-type** {**atm-vcc** | **atm-sdu** | **atm-vpc** | **atm-cell**}] [**vc-switching**]
no apipe service-id

Context config>service

Description The Apipe service provides a point-to-point L2 VPN connection to a remote SAP or to another local SAP. An Apipe can connect an ATM or Frame Relay endpoint either locally or over a PSN to a remote endpoint of the same type or of a different type and perform interworking between the two access technologies.

Parameters *service-id* — The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR on which this service is defined.

Values 1 — 2147483647

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)

vc-type — Specifies a 15 bit value that defines the type of the VC signaled to the peer. Its values are defined in *draft-ietf-pwe3-iana-allocation* and it defines both the signaled VC type as well as the resulting datapath encapsulation over the Apipe.

Values atm-vcc, atm-sdu, atm-vpc, atm-cell

Default atm-sdu

vc-switching — Specifies if the pseudowire switching signalling is used for the spoke SDPs configured in this service.

fpipe

Syntax **fpipe** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*] [**vc-type** {**fr-dlci**}] [**vc-switching**]
no fpipe *service-id*

Context config>service

Description This command configures an Fpipe service. An Fpipe provides a point-to-point L2 VPN connection to a remote SAP or to another local SAP. A Fpipe connects only Frame Relay endpoints either locally or over a PSN to a remote endpoint of the same type.

service-id — The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR on which this service is defined.

Values 1 — 2147483647

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)

vc-type — Specifies a 15 bit value that defines the type of the VC signaled to the peer. Its values are defined in *draft-ietf-pwe3-iana-allocation* and it defines both the signaled VC type as well as the resulting datapath encapsulation over the apipe.

Values fr-dlci

vc-switching — Specifies if the pseudowire switching signalling is used for the spoke SDPs configured in this service.

ipipe

Syntax **ipipe** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**vc-switching**]
no ipipe *service-id*

Context config>service

Description This command configures an IP-Pipe service.

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

Values 1 — 2147483647

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)

vc-switching — Specifies if the pseudowire switching signalling is used for the spoke SDPs configured in this service.

endpoint

Syntax [**no**] **endpoint** *endpoint-name*

Context config>service>apipe
config>service>epipe
config>service>fpipe
config>service>ipipe

Description This command configures a service endpoint.

Parameters *endpoint-name* — Specifies an endpoint name.

active-hold-delay

Syntax **active-hold-delay** *active-hold-delay*
no active-hold-delay

Context config>service>apipe>endpoint
config>service>epipe>endpoint
config>service>fpipe>endpoint
config>service>ipipe>endpoint

Description This command specifies that the node will delay sending the change in the T-LDP status bits for the VLL endpoint when the MC-LAG transitions the LAG subgroup which hosts the SAP for this VLL endpoint from "active" to "standby" or when any object in the endpoint. For example., SAP, ICB, or regular spoke SDP, transitions from up to down operational state.

By default, when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from "active" to "standby", the node sends immediately new T-LDP status bits indicating the new value of "standby" over the spoke SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

There is no delay applied to the VLL endpoint status bit advertisement when the MC-LAG transitions the LAG subgroup which hosts the SAP from "standby" to "active" or when any object in the endpoint transitions to an operationally up state.

Default 0 — A value of zero means that when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from "active" to "standby", the node sends immediately new T-LDP status bits indicating the new value of "standby" over the spoke SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

Parameters *active-hold-delay* — Specifies the active hold delay in 100s of milliseconds.

Values: 0 — 60

revert-time

Syntax **revert-time** [*revert-time* | **infinite**]
no revert-time

Context config>service>apipe>endpoint
config>service>epipe>endpoint
config>service>fpipe>endpoint
config>service>ipipe>endpoint

Description This command configures the time to wait before reverting back to the primary spoke SDP defined on this service endpoint, after having failed over to a backup spoke SDP.

Parameters *revert-time* — Specify the time, in seconds, to wait before reverting to the primary SDP.

Values 0 — 600

infinite — Causes the endpoint to be non-revertive.

interworking

Syntax **interworking {frf-5}**
no interworking

Context config>service>apipe

Description This command specifies the interworking function that should be applied for packets that ingress/egress SAPs that are part of an Apipe service.

Interworking is applicable only when the two endpoints (i.e., the two SAPs or the SAP and the spoke-sdp) are of different types. Also, there are limitations on the combinations of SAP type, vc-type, and interworking values as shown in Table 5.

Table 1: Allowable Combinations of SAP Type, VC-type, and Interworking

SAP Type	Allowed VC-Type Value	Allowed Interworking Value
ATM VC	atm-vcc, atm-sdu	none
	fr-dlci	Not Supported
FR DLCI	fr-dlci	none
	atm-sdu	frf-5

Default **none** (Interworking must be configured before adding a Frame-Relay SAP to an Apipe service.)

Parameters **frf-5 —** Specify Frame Relay to ATM Network Interworking (FRF.5).

service-mtu

Syntax **service-mtu** *octets*
no service-mtu

Context config>service>epipe
config>service>apipe
config>service>fpipe
config>service>ipipe

Description This command configures the service payload (Maximum Transmission Unit – MTU), in bytes, for the service. This MTU value overrides the service-type default MTU.

The **service-mtu** defines the payload capabilities of the service. It is used by the system to validate the SAP and SDP binding's operational state within the service.

The service MTU and a SAP's service delineation encapsulation overhead (i.e., 4 bytes for a dot1q tag) is used to derive the required MTU of the physical port or channel on which the SAP was created. If the required payload is larger than the port or channel MTU, then the SAP will be placed in an inoperative state. If the required MTU is equal to or less than the port or channel MTU, the SAP will be able to transition to the operative state.

When binding an SDP to a service, the service MTU is compared to the path MTU associated with the SDP. The path MTU can be administratively defined in the context of the SDP. The default or administrative path MTU can be dynamically reduced due to the MTU capabilities discovered by the tunneling mechanism of the SDP or the egress interface MTU capabilities based on the next hop in the tunnel path. If the service MTU is larger than the path MTU, the SDP binding for the service will be placed in an inoperative state. If the service MTU is equal to or less than the path MTU, then the SDP binding will be placed in an operational state.

In the event that a service MTU, port or channel MTU, or path MTU is dynamically or administratively modified, then all associated SAP and SDP binding operational states are automatically re-evaluated.

The **no** form of this command returns the default **service-mtu** for the indicated service type to the default value.

Default **apipe, fpipe: 1508**

ipipe: 1500

The following table displays MTU values for specific VC types.

VC-Type	Example Service MTU	Advertised MTU
Ethernet	1514	1500
Ethernet (with preserved dot1q)	1518	1504
VPLS	1514	1500
VPLS (with preserved dot1q)	1518	1504
VLAN (dot1p transparent to MTU value)	1514	1500
VLAN (QinQ with preserved bottom Qtag)	1518	1504

octets — The size of the MTU in octets, expressed as a decimal integer.

Values 1 — 9194

VLL SAP Commands

sap

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

Context config>service>epipe
config>service>apipe
config>service>fpipe
config>service>ipipe

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 7750. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config router interface** *port-type port-id* **mode access** command. Channelized TDM ports are always access ports.

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.

The following are supported:

- ATM VPI/VCI on an ATM port for vc-type atm-vcc and atm-sdu
- ATM VPI on an ATM port for vc-type atm-vpc
- ATM virtual trunk - a range of VPIs on an ATM port for vc-type atm-cell
- ATM port for vc-type atm-cell
- Frame Relay DLCI on a port for vc-type atm-sdu
- ATM SAP carries the IPv4 packet using RFC 2684, VC-Mux or LLC/SNAP routed PDU encapsulation for Ipipe service
- Frame Relay SAP RFC 2427, routed PDU encapsulation for Ipipe service
- Ethernet SAP RFC 1332, PPP IPCP encapsulation of an IPv4 packet for Ipipe service
- Ethernet SAP HDLC SAP uses the routed IPv4 encapsulation for Ipipe service
- ATM - Frame Relay, PPP/IPCP - PPP/IPCP
- Frame Relay-Frame Relay, ATM - ATM
- Ethernet-Ethernet
- cHDLC-cHDLC

- Ethernet SAPs support null, dot1q, and qinq for Ipipe service
- An ATM SAP can be part of a IMA bundle.
- A PPP SAP can be part of a MLPPP bundle.

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 Ethernet Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed.

Default No SAPs are defined.

Special Cases A SAP can be defined with Ethernet ports, SONET/SDH or TDM channels. Up to 32 SAPs can be defined in a service per Media Dependent Adapter (MDA). Attempts to create more than 32 SAPs will generate an error. Up to 49 SDPs can be associated with a service in a single router. Each SDP must have a unique router destination or an error will be generated.

A default SAP has the following format: port-id:*. This type of SAP is supported only on Ethernet MDAs and its creation is allowed only in the scope of Layer 2 services (Epipe and VPLS). This type of SAP is mutually exclusive with a SAP defined by explicit null encapsulation (e.g.m 1/1/1:0).

Two Frame Relay SAPs cannot be configured on an Apipe service. The limitation is for an Apipe service in local mode, which has two SAPs associated with the service, as opposed to a configuration with a SAP and a SDP in remote case, the only combination of the type of SAPs allowed is either two ATM SAPs or an ATM SAP and a Frame Relay SAP. The CLI prevents adding two Frame Relay SAPs under an Apipe service.

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	*slot*/*mda*/*port*[.*channel*]	1/2/3.1
null	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id* | *aps-id*]	*port-id*: 1/2/3 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*: lag-100 *aps-id*: aps-1
dot1q	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id* | *aps-id*]:qtag1	*port-id*:qtag1: 1/2/3:100 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*:qtag1:lag-100:102 *aps-id*:qtag1: aps-1:103
qinq	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	*port-id*:qtag1.qtag2: 1/2/3:100.10 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*:qtag1.qtag2:lag-100:

Type	Syntax	Example
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:vpi/vci |vpi |vpi1.vpi2]	port-id: 9/1/1 aps-id: aps-1 *bundle-id*: bundle-ima-1/1.1 bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 vpi/vci: 16/26 vpi: 16 vpi1.vpi2: 16.200
frame-relay	[*port-id* | *aps-id*]:*dlci*	*port-id*: 1/1/1:100 *aps-id*: aps-1 *dlci*: 16
cisco-hdlc	*slot/mda/port.channel*	*port-id*: 1/2/3.1

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

bundle-id — Specifies the multilink bundle to be associated with this IP interface. The **bundle** keyword must be entered at the beginning of the parameter.
The command syntax must be configured as follows:

bundle-id: **bundle-***type-slot-id*/*mda-slot*.*bundle-num*
bundle-id value range: 1 — 128

For example:

```
*A:ALA-12>config#  port bundle-ppp-5/1.1
*A:ALA-12>config>port# multilink-bundle
```

bgprp-id — Specifies the bundle protection group ID to be associated with this IP interface. The **bpgrp** keyword must be entered at the beginning of the parameter.

The command syntax must be configured as follows:

bpgrp-id: **bpgrp**-*type-bpgrp-num*
type: ima
bpgrp-num value range: 1 — 1280

For example:

```
*A:ALA-12>config# port bpgrp-ima-1
*A:ALA-12>config>service>vpls$ sap bpgrp-ima-1
```

qtag1, qtag2 — Specifies the encapsulation value used to identify the SAP on the port or sub-port. If this parameter is not specificially defined, the default value is 0.

Values qtag1: 0 — 4094
qtag2 : * | 0 — 4094

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 TDM	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 TDM	BCP-Dot1q	0 — 4094	The SAP is identified by the 802.1Q tag on the channel.
SONET/SDH TDM	Frame Relay	16 — 991	The SAP is identified by the data link connection identifier (DLCI).
SONET/SDH ATM	ATM	vpi (NNI) 0 — 4095 vpi (UNI) 0 — 255 vci 1, 2, 5 — 65535 -	The SAP is identified by port or by PVPC or PVCC identifier (vpi, vpi/vci, or vpi range)

endpoint — Adds a SAP endpoint association.

no endpoint — removes the association of a SAP or a spoke-sdp with an explicit endpoint name.

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

agg-rate-limit

Syntax **agg-rate-limit** *agg-rate*
no agg-rate-limit

Context config>service>apipe>sap>egress
config>service>epipe>sap>egress
config>service>fpipe>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, at 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 the queues on the SAP or MSS can operate.

Values 1 — 40000000, max

qinq-mark-top-only

Syntax [**no**] **qinq-mark-top-only**

Context config>service>apipe>sap>egress
config>service>epipe>sap>egress
config>service>fpipe>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 to mark during packet egress. When disabled, both set of P-bits are marked. When the enabled, only the P-bits in the top Q-tag are marked.

Default **no qinq-mark-top-only**

multi-service-site

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

Context config>service>epipe>sap
config>service>apipe>sap
config>service>fpipe>sap
config>service>ipipe>sap

Description This command associates the SAP with a *customer-site-name*. If the specified *customer-site-name* does not exist in the context of the service customer ID an error occurs and the command will not execute. If *customer-site-name* exists, the current and future defined queues on the SAP (ingress and egress) will attempt to use the scheduler hierarchies created within *customer-site-name* as parent schedulers. See multi-service-site on page 82.

The **no** form of the command removes the SAP from any multi-service customer site the SAP belongs to. Removing the site can cause existing or future queues to enter an orphaned state.

Default **None**

customer-site-name — The customer-site-name must exist in the context of the customer-id defined as the service owner. If customer-site-name exists and local scheduler policies have not been applied to the SAP, the current and future queues defined on the SAP will look for their parent schedulers within the scheduler hierarchies defined on customer-site-name.

Values Any valid customer-site-name created within the context of the customer-id.

tod-suite

Syntax **tod-suite** *tod-suite-name*
no tod-suite

Context config>service>apipe>sap
config>service>epipe>sap
config>service>fpipe>sap
config>service>ipipe>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**>**cron** context.

Default **no tod-suite**

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

mac

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

Context config>service>ipipe>sap

Description This command assigns a specific MAC address to an Ipipe SAP.

The **no** form of this command returns the MAC address of the SAP to the default value.

Default The default is the physical MAC address associated with the Ethernet interface where the SAP is configured.

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

mac-refresh

Syntax **mac-refresh** *refresh interval*
no mac-refresh

Context config>service>ipipe>sap

Description This command specifies the interval between ARP requests sent on this Ipipe SAP. When the SAP is first enabled, an ARP request will be sent to the attached CE device and the received MAC address will be used in addressing unicast traffic to the CE. Although this MAC address will not expire while the Ipipe SAP is enabled and operational, it is verified by sending periodic ARP requests at the specified interval.

The **no** form of this command restores mac-refresh to the default value.

Default **14400**

Parameters *refresh interval* — Specifies the interval, in seconds, between ARP requests sent on this Ipipe SAP.

Values 0 — 65535

Service Billing Commands

accounting-policy

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

Context config>service>epipe>sap
config>service>epipe>spoke-sdp
config>service>apipe>sap
config>service>fpipe>sap
config>service>ipipe

Description This command creates the accounting policy context that can be applied to a SAP.

An accounting policy must be defined before it can be associated with a SAP.
If the *policy-id* does not exist, an error message is generated.

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

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

Default Default accounting policy.

Parameters *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 — 99

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>epipe>sap
config>service>epipe>sap
config>service>apipe>sap
config>service>fpipe>sap

Description This command enables accounting and statistical data collection for either the SAP, network port, or IP interface. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOMcards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default **collect-stats**

Service Filter and QoS Policy Commands

egress

Syntax **egress**

Context config>service>epipe>sap
config>service>epipe>spoke-sdp
config>service>apipe>sap
config>service>fpipe>sap
config>service>ipipe>sap

Description This command enables the context to configure egress SAP Quality of Service (QoS) policies.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing.

ingress

Syntax **ingress**

Context config>service>epipe>sap
config>service>epipe>spoke-sdp
config>service>apipe>sap
config>service>fpipe>sap
config>service>ipipe>sap

Description This command enables the context to configure ingress SAP Quality of Service (QoS) policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing.

filter

Syntax **filter** [**ip** *ip-filter-id*]
filter [**ipv6** *ipv6-filter-id*]
filter [**mac** *mac-filter-id*]
no filter [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*] [**mac** *mac-filter-id*]

Context config>service>epipe>sap>egress
config>service>epipe>sap>ingress
config>service>epipe>spoke-sdp>egress
config>service>epipe>spoke-sdp>ingress
config>service>ipipe>spoke-sdp>egress
config>service>ipipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>spoke-sdp>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. Only one filter can be applied to a SAP at a time.

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

IP filters apply only to RFC 2427-routed IP packets. Frames that do not contain IP packets will not be subject to the filter and will always be passed, even if the filter's default action is to drop.

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

Note that IPv6 filters are not supported on a Layer 2 SAP that is configured with QoS MAC criteria. Also, MAC filters are not supported on a Layer 2 SAP that is configured with QoS IPv6 criteria.

Special Cases **Epipe —** Both MAC and IP filters are supported on an Epipe service SAP.

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

Values 1 — 65535

ipv6 *ipv6-filter-id* **—** Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 — 65535

mac *mac-filter-id* **—** Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1 — 65535

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>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>fpipe>spoke-sdp>egress
config>service>fpipe>spoke-sdp>ingress
config>service>ipipe>spoke-sdp>egress
config>service>ipipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>spoke-sdp>ingress

Description This command associates a 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. Only one filter can be applied to a SAP at a time.

The **filter** command is used to associate a filter policy with a specified *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.

IP filters apply only to RFC 2427-routed IP packets. Frames that do not contain IP packets will not be subject to the filter and will always be passed, even if the filter's default action is to drop.

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

ipv6 *ipv6-filter-id* **—** Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 — 65535

qos

Syntax **qos** *policy-id* [**shared-queuing**]
no qos

Context config>service>epipe>sap>egress
config>service>epipe>sap>ingress
config>service>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>sap>ingress

Description This command associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP) or IP interface.

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the *policy-id* does not exist, an error will be returned.

The **qos** command is used to associate both ingress and egress QoS policies. The **qos** command only allows ingress policies to be associated on SAP or IP interface ingress and egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type will return an error.

By default, no specific QoS policy is associated with the SAP or IP interface for ingress or egress, so the default QoS policy is used.

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

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 — This keyword can **only** be specified on SAP ingress. The **shared-queueing** keyword specifies the shared queue policy will be used by this SAP. When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

queue-override

Syntax [**no**] **queue-override**

Context config>service>epipe>sap>egress
config>service>epipe>sap>ingress
config>service>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>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. If the policy was created as a template policy, this command overrides the parameter and its description and queue parameters in the policy.

queue

Syntax [**no**] **queue** *queue-id*

Context config>service>epipe>sap>egress>queue-override
config>service>epipe>sap>ingress>queue-override
config>service>apipe>sap>egress>queue-override
config>service>apipe>sap>ingress>queue-override
config>service>fpipe>sap>egress>queue-override
config>service>fpipe>sap>ingress>queue-override
config>service>ipipe>sap>egress>queue-override
config>service>ipipe>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** *adaptation-rule*}] [**cir** *adaptation-rule*}]
no adaptation-rule

Context config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue
config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>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** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir — The **cir** parameter defines the constraints enforced when adapting the CIR rate defined within the **queue** *queue-id* **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule — Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset.

Values **max** — The **max** (maximum) keyword is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) keyword is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

avg-frame-overhead

Syntax **avg-frame-overhead** *percent*
no avg-frame-overhead

Context config>service>apipe>sap>egress>queue-override>queue
config>service>epipe>sap>egress>queue-override>queue

Description This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a Sonet or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- Offered-load — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- Frame encapsulation overhead — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be 10000 x 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 x 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 x 1.1 or 550 octets.
- Frame based within-cir offered-load — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.

- Frame based PIR — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be 7500 x 1.1 or 8250 octets.
- Frame based within-pir offered-load — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

Port scheduler operation using frame transformed rates — The port scheduler uses the frame based rates to figure 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 0.00 — 100.00

cbs

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

Context config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue
config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue

config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>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 or default

high-prio-only

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

Context config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue
config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>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 execu-

tion. 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*
no mbs

Context config>service>epipe>sap>egress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>apipe>sap>egress>queue-override>queue
config>service>fpipe>sap>egress>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.

Values 0 — 131072 or default

mbs

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

Context config>service>epipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's MBS parameters. The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The 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 returns the MBS size assigned to the queue to the deafult value.

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

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>service>epipe>sap>egress>queue-override>queue
config>service>epipe>sap>ingress>queue-override>queue
config>service>apipe>sap>egress>queue-override>queue
config>service>apipe>sap>ingress>queue-override>queue
config>service>fpipe>sap>egress>queue-override>queue
config>service>fpipe>sap>ingress>queue-override>queue
config>service>ipipe>sap>egress>queue-override>queue
config>service>ipipe>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled 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 — 100000000

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 — 100000000, **max, sum**

Default 0

scheduler-override

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

Context config>service>epipe>sap>egress
config>service>epipe>sap>ingress
config>service>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>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>epipe>sap>egress>sched-override
config>service>epipe>sap>ingress>sched-override
config>service>ipipe>sap>egress>sched-override
config>service>ipipe>sap>ingress>sched-override
config>service>apipe>sap>egress>sched-override
config>service>apipe>sap>ingress>sched-override
config>service>fpipe>sap>egress>sched-override
config>service>fpipe>sap>ingress>sched-override

Description This command can be used to override specific attributes of the specified scheduler name.

A scheduler defines 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>epipe>sap>egress>sched-override>scheduler
config>service>epipe>sap>ingress>sched-override>scheduler
config>service>apipe>sap>egress>sched-override>scheduler
config>service>apipe>sap>ingress>sched-override>scheduler
config>service>fpipe>sap>egress>sched-override>scheduler
config>service>fpipe>sap>ingress>sched-override>scheduler
config>service>ipipe>sap>egress>sched-override>scheduler
config>service>ipipe>sap>ingress>sched-override>scheduler

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 ate 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 queues 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**

scheduler-policy

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

Context config>service>epipe>sap>ingress
config>service>epipe>sap>egress
config>service>apipe>sap>ingress
config>service>apipe>sap>egress
config>service>fpipe>sap>ingress
config>service>fpipe>sap>egress
config>service>ipipe>sap>ingress
config>service>ipipe>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.

match-qinq-dot1p

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

Context config>service>epipe>sap>ingress
config>service>ipipe>sap>ingress

Description This command specifies which Dot1Q tag position 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 6 defines the default behavior for Dot1P evaluation when the **match-qinq-dot1p** command is not executed.

Table 2: 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 7 defines the dot1p evaluation behavior when the top parameter is specified.

Table 3: 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 8 defines the dot1p evaluation behavior when the bottom parameter is specified.

Table 4: 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 4: 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 5: 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 6: 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

Table 6: QinQ Mark Top Only Explicit Marking Actions

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
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 PBits marked using dot1p-value, BottomQ PBits marked with zero
QinQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits PBits marked using dot1p-value, BottomQ PBits marked using preserved value

The QinQ and TopQ SAP PBit marking follows the default behavior devined in Table 9 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.

VLL SDP Commands

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**no-endpoint**]
spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] **endpoint** *endpoint-name* [**icb**]
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>epipe

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 Epipe, VPLS, or VPRN 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 7750 SR devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

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

Special Cases **Epipe —** At most, only one *sdp-id* can be bound to an Epipe service. Since an Epipe is a point-to-point service, it can have, at most, two end points. The two end points can be one SAP and one SDP or two SAPs.

Parameters *sdp-id* — The SDP identifier. Allowed values are integers in the range of 1 to 17407 for existing SDPs.

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

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

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

- The VC type value for Ethernet is 0x0005.
- The VC type value for an Ethernet VLAN is 0x0004.
- The VC type value for a VPLS service is defined as 0x000B.

Values ethernet

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

vlan — Defines the VC type as VLAN. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings. The VLAN VC-type requires at least one dot1Q tag within each encapsulated Ethernet packet transmitted to the far end.

no endpoint — removes the association of a spoke SDP with an explicit endpoint name.

endpoint *endpoint-name* **—** Specifies the name of the service endpoint.

icb — Configures the spoke SDP as an inter-chassis backup SDP binding.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**no-endpoint**]
spoke-sdp *sdp-id*[:*vc-id*] **endpoint** *endpoint-name* [**icb**]
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>apipe
config>service>fpipe
config>service>ipipe

Description This command binds a service to an existing Service Distribution Point (SDP).

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

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a service. If the **sdp** *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 SR/ESS devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

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

Parameters *sdp-id* — The SDP identifier. Allowed values are integers in the range of 1 to 17407 for existing SDPs.

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

no endpoint — Adds or removes a spoke SDP association.

endpoint *endpoint-name* **—** Specifies the name of the service endpoint.

icb — Configures the spoke SDP as an inter-chassis backup SDP binding.

cell-concatenation

Syntax **cell-concatenation**

Context config>service>apipe>spoke-sdp

Description This command enables the context to provide access to the various options that control the termination of ATM cell concatenation into an MPLS frame. Several options can be configured simultaneously. The concatenation process for a given MPLS packet ends when the first concatenation termination condition is met. The concatenation parameters apply only to ATM N:1 cell mode VLL.

aal5-frame-aware

Syntax [**no**] **aal5-frame-aware**

Context config>service>apipe>spoke-sdp>cell-concat

Description This command enables the configuration of AAL5 end-of-message (EOM) to be an indication to complete the cell concatenation operation.

The **no** form of the command resets the configuration to ignore the AAL5 EOM as an indication to complete the cell concatenation.

clp-change

Syntax [**no**] **clp-change**

Context config>service>apipe>spoke-sdp>cell-concat

Description This command enables the configuration of CLP change to be an indication to complete the cell concatenation operation.

The **no** form of the command resets the configuration to ignore the CLP change as an indication to complete the cell concatenation.

max-cells

Syntax **max-cells** *cell-count*
no max-cells [*cell-count*]

Context config>service>apipe>spoke-sdp>cell-concat

Description This command enables the configuration of the maximum number of ATM cells to accumulate into an MPLS packet. The remote peer will also signal the maximum number of concatenated cells it is willing to accept in an MPLS packet. When the lesser of (the configured value and the signaled

value) number of cells is reached, the MPLS packet is queued for transmission onto the pseudowire. It is ensured that the MPLS packet MTU conforms to the configured service MTU.

The **no** form of this command sets max-cells to the value '1' indicating that no concatenation will be performed.

Parameters *cell-count* — Specify the maximum number of ATM cells to be accumulated into an MPLS packet before queueing the packet for transmission onto the pseudowire.

Values 1 — 128

Default 32

max-delay

Syntax **max-delay** *delay-time*
no max-delay [*delay-time*]

Context config>service>apipe>spoke-sdp>cell-concat

Description This command enables the configuration of the maximum amount of time to wait while performing ATM cell concatenation into an MPLS packet before transmitting the MPLS packet. This places an upper bound on the amount of delay introduced by the concatenation process. When this amount of time is reached from when the first ATM cell for this MPLS packet was received, the MPLS packet is queued for transmission onto the pseudowire.

The **no** form of this command resets max-delay to its default value.

Parameters *delay-time* — Specify the maximum amount of time, in hundreds of microseconds, to wait before transmitting the MPLS packet with whatever ATM cells have been received. For example, to bound the delay to 1 ms the user would configure 10 (hundreds of microseconds). The delay-time is rounded up to one of the following values 1, 5, 10, 50, 100, 200, 300 and 400.

Values 1 — 400

control-word

Syntax [**no**] **control-word**

Context config>service>apipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>ipipe>spoke-sdp

Description This command indicates whether the control word is used or not. The value of the control word is negotiated with the peer.

The control word is indicated in the label message as part of pseudowire ID FEC as the C-bit. The control word is mandatory for Apipe service data unit and Fpipe. For other types of Apipe and Epipe the control word is not mandatory. Therefore the usage of the control word is negotiated between the endpoints of the pseudowire.

egress

Syntax **egress**

Context config>service>fpipe>spoke-sdp
config>service>apipe>spoke-sdp
config>service>ipipe>spoke-sdp

Description This command configures the egress SDP context.

ingress

Syntax **ingress**

Context config>service>fpipe>spoke-sdp
config>service>apipe>spoke-sdp

Description This command configures the ingress SDP context.

filter

Syntax **filter** [**ip** *ip-filter-id*]
no filter

Context config>service>fpipe>spoke-sdp>egress
config>service>fpipe>spoke-sdp>ingress

Description This command associates an IP filter policy with an ingress or egress Service Distribution Point (SDP).

Filter policies control the forwarding and dropping of packets based on IP matching criteria. Only one filter can be applied to a spoke SDP at a time.

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 before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

IP filters apply only to RFC 2427-routed IP packets. Frames that do not contain IP packets will not be subject to the filter and will always be passed, even if the filter's default action is to drop.

The **no** form of this command removes any configured filter ID association with the SDP. 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 the **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

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

ip-filter-id — The filter name acts as the ID for the IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 — 65535

vc-label

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

Context config>service>fpipe>spoke-sdp>egress
config>service>apipe>spoke-sdp>egress
config>service>ipipe>spoke-sdp>egress

Description This command configures the egress VC label.

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

Values 16 — 1048575

vc-label

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

Context config>service>fpipe>spoke-sdp>ingress
config>service>apipe>spoke-sdp>ingress
config>service>ipipe>spoke-sdp>ingress

Description This command configures the ingress VC label.

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

Values 2048 — 18431

precedence

Syntax **precedence** [*precedence-value*| **primary**]
no precedence

Context config>service>apipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>ipipe>spoke-sdp

Description This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding will begin to forward traffic.

The no form of the command returns the precedence value to the default.

Default **4**

Parameters *precedence-value* — Specifies the spoke SDP precedence.

Values 1 — 4

primary — Specifies to make this the primary spoke SDP.

vc-label

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

Context config>service>epipe>spoke-sdp>egress

Description This command configures the egress VC label.

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

Values 16 — 1048575

vc-label

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

Context config>service>epipe>spoke-sdp>ingress

Description This command configures the ingress VC label.

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

Values 2048 — 18431

vlan-vc-tag

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

Context config>service>epipe>spoke-sdp

Description This command specifies an explicit dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured dot1q tag can be overridden by a received TLV specifying the dot1q value expected by the far end. This signaled value must be stored as the remote signaled dot1q value for the binding. The provisioned local dot1q tag must be stored as the administrative dot1q value for the binding.

When the dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.

The **no** form of this command disables the command

Default **no vlan-vc-tag**

Parameters *0..4094* — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

ATM Commands

atm

Syntax **atm**

Context config>service>epipe>sap
config>service>apipe>sap
config>service>ipipe>sap
config>service>epipe>sap

Description This command enables access to the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as:

- Configuring ATM port or ATM port-related functionality on MDAs supporting ATM functionality
- Configuring ATM-related configuration for ATM-based SAPs that exist on MDAs supporting ATM functionality.

If ATM functionality is not supported for a given context, the command returns an error.

egress

Syntax **egress**

Context config>service>epipe>sap
config>service>epipe>sap>atm
config>service>apipe>sap>atm
config>service>fpipe>sap

This command configures egress ATM attributes for the SAP.

ingress

Syntax **ingress**

Context config>service>epipe>sap
config>service>epipe>sap>atm
config>service>epipe>sap
config>service>apipe>sap>atm

Description This command configures ingress ATM attributes for the SAP.

encapsulation

Syntax **encapsulation** *atm-encap-type*

Context config>service>epipe>sap>atm
config>service>ipipe>sap>atm

Description This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition references RFC 2684, *Multiprotocol Encapsulation over ATM AAL5*, and to the ATM Forum LAN Emulation specification.

Ingress traffic that does not match the configured encapsulation will be dropped.

Default The encapsulation is driven by the services for which the SAP is configured.
For IES and VPRN service SAPs, the default is **aal5snap-routed**.

Parameters *atm-encap-type* — Specify the encapsulation type.

Values **aal5snap-routed** — Routed encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684.
aal5mux-ip — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684

traffic-desc

Syntax **traffic-desc** *traffic-desc-profile-id*
no traffic-desc

Context config>service>epipe>sap
config>service>apipe>sap>atm>egress
config>service>apipe>sap>atm>ingress
config>service>epipe>sap>atm>egress
config>service>epipe>sap>atm>ingress

Description This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP).

When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction.

When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.

The **no** form of the command reverts the traffic descriptor to the default traffic descriptor profile.

Default The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.

Parameters *traffic-desc-profile-id* — Specify a defined traffic descriptor profile (see the QoS atm-td-profile command).

oam

Syntax **oam**

Context config>service>epipe>sap
config>service>apipe>sap>atm

Description This command enables the context to configure OAM functionality for a PVCC delimiting a SAP.

- The ATM-capable MDAs support end-to-end and segment OAM functionality (AIS, RDI, Loop-back) over both F5 (VC) and end-to-end F4 (VP) OAM:
- ITU-T Recommendation I.610 - B-ISDN Operation and Maintenan version 11/95
- GR-1248-CORE - Generic Requirements for Operations of ATM N3 June 1996
- GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) (AAL) Protocols Generic Requirements, Issue 1, July 1994

alarm-cells

Syntax [**no**] **alarm-cells**

Context config>service>epipe>sap>oam
config>service>epipe>sap>oam
config>service>apipe>sap>atm>oam

Description This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC terminations to monitor and report the status of their connection by propagating fault information through the network and by driving PVCC's operational status.

When alarm-cells functionality is enabled, a PVCC's operational status is affected when a PVCC goes into an AIS or RDI state because of an AIS/RDI processing (assuming nothing else affects PVCC's operational status, for example, if the PVCC goes DOWN, or enters a fault state and comes back UP, or exits that fault state). RDI cells are generated when PVCC is operationally DOWN. No OAM-specific SNMP trap is raised whenever an endpoint enters/exits an AIS or RDI state, however, if as result of an OAM state change, the PVCC changes operational status, then a trap is expected from an entity the PVCC is associated with (for example a SAP).

The **no** command disables alarm-cells functionality for a PVCC. When alarm-cells functionality is disabled, a PVCC's operational status is no longer affected by a PVCC's OAM state changes due to AIS/RDI processing (Note that when alarm-cells is disabled, a PVCC will change operational status to UP due to alarm-cell processing) and RDI cells are not generated as result of the PVCC going into AIS or RDI state. The PVCC's OAM status, however, will record OAM faults as described above.

Default Enabled for PVCCs delimiting IES SAPs

ce-address

Syntax **ce-address** *ip-address*
no ce-address

Context config>service>ipipe>sap

config>service>ipipe>spoke-sdp

Description This command specifies the IP address of the CE device associated with an Ipipe SAP or spoke SDP. In the case of a SAP, it is the address of the CE device directly attached to the SAP. For a spoke SDP, it is the address of the CE device reachable through that spoke SDP (for example, attached to the SAP on the remote node). The address must be a host address (no subnet addresses are accepted) as there must be only one CE device attached to an Ipipe SAP. The CE address specified at one end of an Ipipe will be used in processing ARP messages at the other endpoint, as the router acts as a proxy for ARP messages.

Default **none**

terminate

Syntax [**no**] **terminate**

Context config>service>apipe>sap>atm>oam

Description This command specifies whether this SAP will act as an OAM termination point. ATM SAPs can be configured to tunnel or terminate OAM cells.

When configured to not terminate (the default is **no terminate**), the SAP will pass OAM cells through the VLL without inspecting them. The SAP will respond to OAM loopback requests that are directed to the local node by transmitting a loopback reply. Other loopback requests are transparently tunneled through the PW. In this mode, it is possible to launch a loopback request towards the directly-attached ATM equipment and see the results of the reply.

When configured to terminate, the SAP will respond to AIS by transmitting RDI and will signal the change of operational status to the other endpoint (e.g., through LDP status notifications). The SAP will respond to OAM loopback requests by transmitting a loopback reply. In this mode, it is possible to launch a loopback request towards the directly-attached ATM equipment and see the results of the reply. This option is available only for ATM-SDU pseudowire connections, and it is not supported for VT and port pseudowire connections.

Default **no terminate**

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.

Table 7: 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 2049 — 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** [**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 Displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

Parameters **ingress —** Specifies matching an ingress policy.

egress — Specifies matching an egress policy.

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

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]

	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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.

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.
Sap MTU	The SAP MTU value.
I.QoS	The SAP ingress QoS policy number specified on the ingress SAP.
I.MAC/IP	The MAC or IP filter policy ID applied to the ingress SAP.
E.QoS	The SAP egress QoS policy number specified on the egress SAP.
E.Mac/IP	The MAC or IP filter policy ID applied to the egress SAP
A.Pol	The accounting policy ID assigned to the SAP.
Adm	The desired state of the SAP.
Opr	The actual state of the SAP.

Sample Output

```
*A:ALA-48# show service sap-using
===============================================================================
Service Access Points
===============================================================================
PortId                 SvcId      Ing.  Ing.    Egr.  Egr.    Anti   Adm  Opr
                                  QoS   Fltr    QoS   Fltr    Spoof
```

```
-------------------------------------------------------------------------------
1/1/21:0              1          1     none    1     none    none   Up   Down
1/1/4:0               3          2     none    1010  none    n/a    Up   Up
9/2/1:0/32            5          2     none    1010  none    n/a    Up   Down
1/2/9:0               6          1     none    1     none    n/a    Up   Down
1/1/6:0               7          1     ip4     1     none    n/a    Up   Down
2/1/8:0               7          1     none    1     ip4     n/a    Up   Down
8/1/2:4094            8          1     none    1     none    n/a    Up   Down
1/2/10:100            12         1     none    1     none    n/a    Up   Down
2/1/10:0              13         1     none    1     none    none   Up   Down
2/2/2:200             14         1     none    1     none    n/a    Up   Down
1/2/19:0              88         2     ip4     1     none    ip-mac Up   Down
1/2/20:0              88         4     ip4     1     ip4     none   Up   Down
3/2/4:50/5            88         1     none    1     none    none   Up   Down
1/1/18:0              89         1     none    1     none    none   Up   Down
1/1/7:0               103        1     none    1     none    none   Up   Down
1/1/11:0              103        1     none    1     none    none   Up   Down
1/1/2:445             202        1     none    1     none    n/a    Up   Down
1/1/2:555             202        1     none    1     none    n/a    Up   Down
1/1/2:446             204        1     none    1     none    n/a    Up   Up
2/2/2:16              204        1     none    1     none    n/a    Up   Down
1/1/2:447             206        1     none    1     none    n/a    Up   Up
1/1/2:448             206        1     none    1     none    n/a    Up   Up
1/1/9:0               700        100   none    1     ip4     none   Up   Down
2/1/4:100             800        1     none    1     none    none   Up   Down
2/1/4:200             800        1     none    1     none    none   Up   Down
2/1/5:16              9000       1     none    1     none    none   Up   Down
8/1/3:0               32806      1     none    1     none    none   Up   Down
8/1/3:22              32806      1     none    1     none    none   Up   Down
1/2/2:0               90001      1     none    1     none    none   Up   Down
2/1/5:0               90001      1     none    1     none    none   Up   Down
1/2/14:0              10203041   1     none    1     none    none   Up   Down
-------------------------------------------------------------------------------
Number of SAPs : 31
-------------------------------------------------------------------------------
===============================================================================
*A:ALA-48#
*A:ALA-48# show service sap-using sap 1/1/21:0

===============================================================================
Service Access Points Using Port 1/1/21:0
===============================================================================
PortId                SvcId      Ing.  Ing.    Egr.  Egr.    Anti   Adm  Opr
                                 QoS   Fltr    QoS   Fltr    Spoof
-------------------------------------------------------------------------------
1/1/21:0              1          1     none    1     none    none   Up   Down
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALA-48#

*A:ALA-12#  show service sap-using egress atm-td-profile 2
===============================================================================
Service Access Point Using ATM Traffic Profile 2
===============================================================================
PortId     SvcId    I.QoS I.Fltr E.QoS E.Fltr A.Pol Adm Opr
-------------------------------------------------------------------------------
5/1/1:0/11 511111    2    none     2    none   none  Up  Up
5/1/1:0/12 511112    2    none     2    none   none  Up  Up
5/1/1:0/13 511113    2    none     2    none   none  Up  Up
5/1/1:0/14 511114    2    none     2    none   none  Up  Up
5/1/1:0/15 511115    2    none     2    none   none  Up  Up
```

```
5/1/1:0/16 511116    2    none     2    none    none  Up   Up
5/1/1:0/17 511117    2    none     2    none    none  Up   Up
5/1/1:0/18 511118    2    none     2    none    none  Up   Up
5/1/1:0/19 511119    2    none     2    none    none  Up   Up
5/1/1:0/20 511120    2    none     2    none    none  Up   Up
5/1/1:0/21 511121    2    none     2    none    none  Up   Up
5/1/1:0/22 511122    2    none     2    none    none  Up   Up
5/1/1:0/23 511123    2    none     2    none    none  Up   Up
5/1/1:0/24 511124    2    none     2    none    none  Up   Up
5/1/1:0/25 511125    2    none     2    none    none  Up   Up
...
===============================================================================
*A:ALA-12#
```

sdp

Syntax **sdp** [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]

Context show>service

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 desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Opr MTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
IP address	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.

Label	Description
Adm Admin State	Specifies the desired state of the SDP.
Opr Oper State	Specifies the operating state of the SDP.
Deliver Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
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.
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. timer expired.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
TX Hello Msgs	Specifies the number of SDP echo request messages transmitted since the keepalive was administratively enabled or the counter was cleared.
Rx Hello Msgs	Specifies the number of SDP echo request messages received since the keepalive was administratively enabled or the counter was cleared.

Label	Description
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
===============================================================================
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:ALA-12# show service sdp 8
===============================================================================
Service Destination Point (Sdp Id : 8)
===============================================================================
SdpId    Adm MTU   Opr MTU   IP address       Adm  Opr         Deliver Signal
-------------------------------------------------------------------------------
8        4462      4462      10.10.10.104     Up   Dn NotReady MPLS    TLDP
===============================================================================
*A:ALA-12#
===============================================================================
Service Destination Point (Sdp Id : 8) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 8  -(10.10.10.104)
-------------------------------------------------------------------------------
Description           : MPLS-10.10.10.104
SDP Id               : 8
Admin Path MTU       : 0                    Oper Path MTU      : 0
Far End              : 10.10.10.104         Delivery           : MPLS
Admin State          : Up                   Oper State         : Down
Flags                : SignalingSessDown TransportTunnDown
Signaling            : TLDP                 VLAN VC Etype      : 0x8100
Last Status Change   : 02/01/2007 09:11:39  Adv. MTU Over.     : No
Last Mgmt Change     : 02/01/2007 09:11:46

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

Associated LSP LIST :
Lsp Name             : to-104
Admin State          : Up                   Oper State         : Down
Time Since Last Tran*: 01d07h36m
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:ALA-12#
```

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.
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#
```

service-using

Syntax **service-using** [**epipe**] [**ies**] [**vpls**] [**vprn**] [**mirror**] [**apipe**] [**fpipe**] [**ipipe**] [**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 [**service**] — Displays specified service information.

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 any 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
-------------------------------------------------------------------------------
Matching Services : 3
-------------------------------------------------------------------------------
===============================================================================
*A:ALA-12#


*A:ALA-12# show service service-using epipe
===============================================================================
Services [epipe]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
6            Epipe     Up     Up         6                 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:ALA-12# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
1            uVPLS     Up     Up         1                 10/26/2006 15:44:57
2            Epipe     Up     Down       1                 10/26/2006 15:44:57
10           mVPLS     Down   Down       1                 10/26/2006 15:44:57
11           mVPLS     Down   Down       1                 10/26/2006 15:44:57
100          mVPLS     Up     Up         1                 10/26/2006 15:44:57
101          mVPLS     Up     Up         1                 10/26/2006 15:44:57
102          mVPLS     Up     Up         1                 10/26/2006 15:44:57
999          uVPLS     Down   Down       1                 10/26/2006 16:14:33
-------------------------------------------------------------------------------
Matching Services : 8
-------------------------------------------------------------------------------
*A:ALA-12#
```

id

Syntax **id** *service-id* {**all** | **arp** | **base** | **endpoint** | **fdb** | **interface** | **labels** | **sap** | **sdp** | **split-horizon-group** | **stp**}

Context show>service

Description Display information for a particular service-id.

Parameters *service-id* — The service identification number that identifies the service in the domain.

all — Display detailed information about the service.

arp — Display ARP entries for the service.

base — Display basic service information.

endpoint — Display service endpoint 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.

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 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.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255

vci	1, 2, 5 — 65535
dlci	16 — 1022

Sample Output

```
*A:ALA-1# show service id 11 authentication statistics
===============================================================
Authentication statistics
===============================================================
Interface / SAP                Authentication  Authentication
                               Successful      Failed
---------------------------------------------------------------
vpls-11-90.1.0.254              1582            3
---------------------------------------------------------------
Number of entries: 1
===============================================================
*A:ALA-1#
```

all

Syntax **all**

Context show>service>id

Description 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.
VLL Type	Specifies the VLL type.
SDP Id	The SDP identifier.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
SAP Count	The number of SAPs specified for this service.
SDP Bind Count	The number of SDPs bound to this service.
Split Horizon Group specifics	

Label	Description (Continued)
Split Horizon Group	Name of the split horizon group for this VPLS.
Description	Description of the split horizon group.
Last Changed	The date and time of the most recent management-initiated change to this split horizon group.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.
Type	Indicates whether this Service SDP binding is a spoke or a mesh.
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: 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.
Keepalive	
Admin State	Specifies the admin. state of the keepalive protocol.
Oper State	Specifies the operational state of the keepalive protocol.
Hello Time	Specifies how often the SDP Echo Request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP Echo Request messages transmitted on this SDP.

Label	Description (Continued)
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. 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 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 of.
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.

Label	Description (Continued)
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
SAP Statistics	
Dropped	The number of packets or octets dropped.
Offered Hi Priority	The number of high priority packets, as determined by the SAP ingress QoS policy.
Offered Low Priority	The number of low priority packets, as determined by the SAP ingress QoS policy.
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
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.

Label	Description (Continued)
In profile dropped	The number of in-profile packets or octets dropped for the SAP.
Out profile forwarded	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Out profile dropped	The number of out-of-profile packets or octets discarded.
DHCP Relay	
State	Specifies whether DHCP Relay is enabled on this SAP.
Info Option	Specifies whether Option 82 processing is enabled on this SAP.
Action	Specifies the Option 82 processing on this SAP or interface: keep, replace or drop.
Circuit ID	Specifies whether the If Index is inserted in Circuit ID suboption of Option 82.
Remote ID	Specifies whether the far-end MAC address is inserted in Remote ID suboption of Option 82
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.

Sample Output

```
*B:ALA-Dut-H# show service id 100 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 100                Vpn Id            : 100
Service Type      : Epipe
Description       : Default epipe description for service id 100
Customer Id       : 100
Last Status Change: 02/06/2007 10:03:11
Last Mgmt Change  : 02/06/2007 09:43:27
Admin State       : Up                 Oper State        : Up
MTU               : 1514
```

```
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 4
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 1:10100  -(10.20.1.7)
-------------------------------------------------------------------------------
SDP Id             : 1:10100                 Type              : Spoke
VC Type            : Ether                   VC Tag            : n/a
Admin Path MTU     : 1560                    Oper Path MTU     : 1560
Far End            : 10.20.1.7               Delivery          : MPLS

Admin State        : Up                      Oper State        : Up
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : 130065                  Egress Label      : 130368
Ing mac Fltr       : n/a                     Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                     Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                     Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred           Oper ControlWord  : False
Last Status Change : 02/06/2007 10:03:24     Signaling         : TLDP
Last Mgmt Change   : 02/06/2007 09:43:27
Endpoint           : y                       Precedence        : 4
Flags              : SapOperDown
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : mplsRouterAlertLabel
MAC Pinning        : Disabled

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

Statistics           :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0

Associated LSP LIST :
Lsp Name           : lsp1_G
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 01h40m15s

-------------------------------------------------------------------------------
 Sdp Id 2:100  -(10.20.1.4)
-------------------------------------------------------------------------------
SDP Id             : 2:100                   Type              : Spoke
VC Type            : Ether                   VC Tag            : n/a
Admin Path MTU     : 1560                    Oper Path MTU     : 1560
Far End            : 10.20.1.4               Delivery          : MPLS

Admin State        : Up                      Oper State        : Up
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : 130671                  Egress Label      : 130367
Ing mac Fltr       : n/a                     Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                     Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                     Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred           Oper ControlWord  : False
Last Status Change : 02/06/2007 10:03:11     Signaling         : TLDP
Last Mgmt Change   : 02/06/2007 09:43:27
Endpoint           : y                       Precedence        : 0
Flags              : None
```

```
Peer Pw Bits       : pwFwdingStandby
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : mplsRouterAlertLabel
MAC Pinning        : Disabled

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

Statistics           :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0

Associated LSP LIST :
Lsp Name           : lsp2_D
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 01h40m16s

-------------------------------------------------------------------------------
 Sdp Id 3:100  -(10.20.1.5)
-------------------------------------------------------------------------------
SDP Id             : 3:100                   Type              : Spoke
VC Type            : Ether                   VC Tag            : n/a
Admin Path MTU     : 1560                    Oper Path MTU     : 1560
Far End            : 10.20.1.5               Delivery          : MPLS

Admin State        : Up                      Oper State        : Up
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : 130971                  Egress Label      : 130368
Ing mac Fltr       : n/a                     Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                     Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                     Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred           Oper ControlWord  : False
Last Status Change : 02/06/2007 10:03:17     Signaling         : TLDP
Last Mgmt Change   : 02/06/2007 09:43:27
Endpoint           : y                       Precedence        : 4
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : mplsRouterAlertLabel
MAC Pinning        : Disabled

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

Statistics           :
I. Fwd. Pkts.      : 0                       I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0

Associated LSP LIST :
Lsp Name           : lsp3_E
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 01h40m16s

-------------------------------------------------------------------------------
...
===============================================================================
*B:ALA-Dut-H#
```

```
*A:ALA-DutC>config>service# show service id 100 all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 100                 Vpn Id            : 1
Service Type      : Apipe               VLL Type          : AAL5SDU
Description       : Default apipe description for service id 100
Customer Id       : 1
Last Status Change: 04/04/2007 20:53:13
Last Mgmt Change  : 04/04/2007 20:48:24
Admin State       : Up                  Oper State        : Up
MTU               : 1508
Vc Switching      : True
SAP Count         : 0                   SDP Bind Count    : 2

-------------------------------------------------------------------------------
APIPE SDU-mode specifics
-------------------------------------------------------------------------------
Interworking      : None

-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 3:1  -(1.1.1.1)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id            : 3:1                       Type              : Spoke
VC Type           : AAL5SDU                   VC Tag            : 0
Admin Path MTU    : 1600                      Oper Path MTU     : 1600
Far End           : 1.1.1.1                   Delivery          : GRE

Admin State       : Up                        Oper State        : Up
Acct. Pol         : None                      Collect Stats     : Disabled
Ingress Label     : 119665                    Egress Label      : 103665
Ing mac Fltr      : n/a                       Egr mac Fltr      : n/a
Ing ip Fltr       : n/a                       Egr ip Fltr       : n/a
Admin ControlWord : Preferred                 Oper ControlWord  : True
Last Status Change : 04/04/2007 20:52:24      Signaling         : TLDP
Last Mgmt Change  : 04/04/2007 20:48:24
Endpoint          : N/A                       Precedence        : 4
Flags             : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : pwe3ControlWord
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
E. Fwd. Pkts.     : 0                         E. Fwd. Octets    : 0

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
```

```
-------------------------------------------------------------------------------
 Sdp Id 6:2  -(4.4.4.4)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id            : 6:2                       Type              : Spoke
VC Type           : AAL5SDU                   VC Tag            : 0
Admin Path MTU    : 1600                      Oper Path MTU     : 1600
Far End           : 4.4.4.4                   Delivery          : GRE

Admin State       : Up                        Oper State        : Up
Acct. Pol         : None                      Collect Stats     : Disabled
Ingress Label     : 103664                    Egress Label      : 119665
Ing mac Fltr      : n/a                       Egr mac Fltr      : n/a
Ing ip Fltr       : n/a                       Egr ip Fltr       : n/a
Admin ControlWord : Preferred                 Oper ControlWord  : True
Last Status Change : 04/04/2007 20:53:13      Signaling         : TLDP
Last Mgmt Change  : 04/04/2007 20:48:24
Endpoint          : N/A                       Precedence        : 4
Flags             : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : pwe3ControlWord mplsRouterAlertLabel
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
E. Fwd. Pkts.     : 0                         E. Fwd. Octets    : 0

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS

-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
No Sap Associations

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:ALA-DutC>config>service#
```

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	The type of service: Epipe, Apipe, Fpipe, Ipipe, VPLS, IES, VPRN.
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:ALA-12# show service id 1 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPRN
Customer Id       : 1
Last Status Change: 02/01/2007 09:11:39
Last Mgmt Change  : 02/01/2007 09:11:46
```

```
Admin State       : Up                 Oper State        : Down

Route Dist.       : 10001:1
AS Number         : 10000              Router Id         : 10.10.10.103
ECMP              : Enabled            ECMP Max Routes   : 8
Max Routes        : No Limit           Auto Bind         : LDP
Vrf Target        : target:10001:1
Vrf Import        : vrfImpPolCust1
Vrf Export        : vrfExpPolCust1
SAP Count         : 1                  SDP Bind Count    : 18

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                      Type         AdmMTU  OprMTU  Adm     Opr
-------------------------------------------------------------------------------
sap:1/1/7:0                     q-tag        1518    1518    Up      Up
sdp:10:1 M(10.20.1.3)           TLDP         4462    4462    Up      TLDP Down
sdp:20:1 M(10.20.1.4)           TLDP         4462    4462    Up      TLDP Down
sdp:30:1 M(10.20.1.5)           TLDP         4462    4462    Up      TLDP Down
sdp:40:1 M(10.20.1.20)          TLDP         1534    4462    Up      Up
sdp:200:1 M(10.20.1.30)         TLDP         1514    4462    Up      Up
sdp:300:1 M(10.20.1.31)         TLDP         4462    4462    Up      TLDP Down
sdp:500:1 M(10.20.1.50)         TLDP         4462    4462    Up      TLDP Down
===============================================================================
*A:ALA-12#
```

endpoint

Syntax **endpoint** [*endpoint-name*]

Context show>service>id

Description This command displays service endpoint information.

Parameters *endpoint-name* — Specifies the name of an existing endpoint for the service.

Sample Output

```
*A:ALA-48>config>service>epipe# show service id 6 endpoint
===============================================================================
Service 6 endpoints
===============================================================================
Endpoint name      : x
Revert time        : 0
Act Hold Delay     : 0
Tx Active          : none
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
No members found.
===============================================================================
Endpoint name      : y
Revert time        : 0
Act Hold Delay     : 0
Tx Active          : none
-------------------------------------------------------------------------------
```

```
Members
------------------------------------------------------------------------------
No members found.
==============================================================================
==============================================================================
*A:ALA-48>config>service>epipe#
```

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#
```

lease-state

Syntax **lease-state** [[**sap** *sap-id*] [**sdp** [*sdp-id*[:*vc-id*]]] | [**interface** *interface-name*] | [**ip-address** *ip-address[/mask*]] | [**mac** *ieee-address*] | [**wholesaler** *service-id*]] [**detail**]

Context show>service>id>dhcp

Description This command displays DHCP lease state information. Note that the **wholesaler** *service-id* parameter is applicable only in the VPRN context.

Parameters **interface *interface-name* —** Displays information for the specified IP interface.

ip-address *ip-address* — Displays information associated with the specified IP address.

detail — Displays detailed information.

wholesaler *service-id* **—** The VPRN service ID of the wholesaler. When specified in this context, SAP, SDP, interface, IP address and MAC parameters are ignored.

Values 1 — 2147483647

Values

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	

vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

Sample Output

```
*A:ALA-007# show service id 1000 dhcp lease-state
===============================================================================
DHCP lease state table, service 1000
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining Lease    MC
                                                      LifeTime Origin   Stdby
-------------------------------------------------------------------------------
13.13.40.254    00:00:00:00:03:13 1/1/1:13            00h04m54s
DHCP-R
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:ALA-007#

*A:ALA-007#show service id 1000 dhcp lease-state detail
===============================================================================
DHCP lease states for service 1000
===============================================================================
Service ID           : 1000
IP Address           : 13.13.40.254
Mac Address          : 00:00:00:00:03:13
Subscriber-interface : whole-sub
Group-interface      : intf-13
Retailer             : 2000
Retailer If          : retail-sub
SAP                  : 1/1/1:13
Remaining Lifetime   : 00h04m50s (Lease Split)
Persistence Key      : N/A
Sub-Ident            : "Alcatel-Lucent"
Sub-Profile-String   : "ADSL GO"
SLA-Profile-String   : "BE-Video"
Lease ANCP-String    : ""
Sub-Ident origin     : Retail Radius
Strings origin       : Retail Radius
Lease Info origin    : Retail DHCP
Ip-Netmask           : 255.255.0.0
Broadcast-Ip-Addr    : 13.13.255.255
Default-Router       : N/A
Primary-Dns          : N/A
Secondary-Dns        : N/A

ServerLeaseStart     : 04/24/2003 13:35:38
ServerLastRenew      : 04/24/2003 13:35:38
ServerLeaseEnd       : 04/24/2003 15:15:38
Session-Timeout      : 0d 01:40:00
DHCP Server Addr     : 10.232.237.2

Persistent Relay Agent Information
  Circuit Id         : 1/1/1:13
  Remote Id          : stringtest
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:ALA-007
```

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.

wholesalers

Syntax **wholesalers**

Context show>service>id

Description This command displays service wholesaler information.

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 form *slot*/*mda*/*port*[**.***channel*].

interface *interface-name* — Displays information for the specified IP interface.

ip-address *ip-address* — Displays information associated with the specified IP address.

detail — Displays detailed information.

wholesaler *service-id* **—** The VPRN service ID of the wholesaler. When specified in this context, SAP, SDP, interface, IP address and MAC parameters are ignored.

Values 1 — 2147483647

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 SAP and qtag.
Encap	The encapsulation type of the SAP.
Ethertype	Specifies an Ethernet type II Ethertype value.

Label	Description
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.
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.
Forwarding Engine Stats	
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.
Off.Uncolor	The number of uncolored packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Queueing Stats (Ingress QoS Policy)	
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.

Label	Description
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.
Queuing Stats (Egress QoS Policy)	
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 Han-dling	The indication that OAM cells are being processed.
AAL-5 Encap	The AAL-5 encapsulation type.

Sample Output

```
A:ALA-48>config>service>epipe# show service id 8 sap 8/1/2:4094
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 8
SAP                : 8/1/2:4094              Encap             : bcpDot1q

Admin State        : Up                      Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown
Last Status Change : 02/06/2007 12:01:14
Last Mgmt Change   : 02/06/2007 12:01:17
Admin MTU          : 1522                    Oper MTU          : 1522
Ingress qos-policy : 1                       Egress qos-policy : 1
Shared Q plcy      : n/a                     Multipoint shared : Disabled
Ingress Filter-Id  : n/a                     Egress Filter-Id  : n/a
tod-suite          : None

Multi Svc Site     : None
Acct. Pol          : None                    Collect Stats     : Disabled
===============================================================================
A:ALA-48>config>service>epipe#


A:ALA-48>config>service>epipe# show service id 8 sap 8/1/2:4094 detail
```

```
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 8
SAP                : 8/1/2:4094              Encap             : bcpDot1q

Admin State        : Up                      Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown
Last Status Change : 02/06/2007 12:01:14
Last Mgmt Change   : 02/06/2007 12:01:17
Admin MTU          : 1522                    Oper MTU          : 1522
Ingress qos-policy : 1                       Egress qos-policy : 1
Shared Q plcy      : n/a                     Multipoint shared : Disabled
Ingress Filter-Id  : n/a                     Egress Filter-Id  : n/a
tod-suite          : None

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 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
===============================================================================
A:ALA-48>config>service>epipe#
```

If a TOD Suite is configured on a SAP, the name of the suite is shown in the show command output. The values of the policies may be different from those configured on the SAP, because the configured policy assignments may have been overruled by policy assignments of the TOD Suite.

Sample Output

```
A:ALA-48# show service id 1 sap 1/1/1:2
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/1/1:5                  Encap             : q-tag
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100

Admin State        : Up                       Oper State        : Up
Flags              : None
Last Status Change : 10/05/2006 17:06:03
Last Mgmt Change   : 10/05/2006 22:30:03
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
Admin MTU          : 1518                     Oper MTU          : 1518
Ingress qos-policy : 1190                     Egress qos-policy : 1190
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          : suite_sixteen            qinq-pbit-marking : both
Egr Agg Rate Limit : max
ARP Reply Agent    : Unknown                  Host Conn Verify  : Disabled
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
L2PT Termination   : Disabled                 BPDU Translation  : Disabled

Multi Svc Site     : None
I. Sched Pol       : SchedPolCust1_Night
E. Sched Pol       : SchedPolCust1Egress_Night
Acct. Pol          : None                     Collect Stats     : Disabled

Anti Spoofing      : None                     Nbr Static Hosts  : 0
===============================================================================
A:ALA-48#

A:kerckhot_4# show service id 1 sap 1/1/1:6
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/1/1:6                  Encap             : q-tag
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100

Admin State        : Up                       Oper State        : Down
Flags              : TodResourceUnavail
Last Status Change : 12/01/2006 09:59:42
Last Mgmt Change   : 12/01/2006 09:59:45
...
A:kerckhot_4#
```

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 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, vlan, or vpls.
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 current state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.

Label	Description (Continued)
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the Keepalive process.
Oper State	The operational state of the Keepalive process.
Hello Time	Specifies how often the SDP Echo Request messages are transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	Specifies the length of the SDP Echo Request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
E. Fwd. Octets	Specifies the number of forwarded egress octets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. If the SDP type is GRE, then the following message displays: SDP Delivery Mechanism is not MPLS

Sample Output

```
A:SR1_3# show service id 2 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 900:3  -(1.1.1.1)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id            : 900:3                   Type             : Spoke
VC Type           : AAL5SDU                 VC Tag           : 0
Admin Path MTU    : 1522                    Oper Path MTU    : 1522
Far End           : 1.1.1.1                 Delivery         : LDP

Admin State       : Up                      Oper State       : Up
Acct. Pol         : None                    Collect Stats    : Disabled
Ingress Label     : 131069                  Egress Label     : 131069
Ingress Filter    : n/a                     Egress Filter    : n/a
```

```
Last Status Change : 10/04/2006 16:12:15      Signaling         : TLDP
Last Mgmt Change   : 10/04/2006 16:10:50
Flags              : None
Peer Pw Bits       : lacIngressFault lacEgressFault psnEgressFault
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : mplsRouterAlertLabel
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.      : 2                        I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 13                       E. Fwd. Octets    : 676
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
A:SR1_3#


A:ALA-49# show service id 6 sdp 2:6 detail
===============================================================================
Service Destination Point (Sdp Id : 2:6) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 2:6  -(10.10.10.103)
-------------------------------------------------------------------------------
Description     : GRE-10.10.10.103
SDP Id             : 2:6                      Type              : Spoke
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 1472
Far End            : 10.10.10.103             Delivery          : GRE

Admin State        : Up                       Oper State        : Down
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 6298                     Egress Label      : 6300
Ingress Filter     : n/a                      Egress Filter     : n/a
Last Status Change : 10/07/2006 12:30:02      Signaling         : TLDP
Last Mgmt Change   : 10/07/2006 12:30:03
Flags              : SapOperDown
                     PathMTUTooSmall
Peer Pw Bits       : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
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

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS

-------------------------------------------------------------------------------
```

```
Number of SDPs : 1
------------------------------------------------------------------------------
==============================================================================
A:ALA-49#
```

control-word

The following examples show both sides (PE nodes) when control word is enabled:

```
*A:ALA-Dut-B>config>service>epipe# show service id 2100 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 1:2001  -(1.1.1.1)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id            : 1:2001                  Type              : Spoke
VC Type           : Ether                   VC Tag            : n/a
Admin Path MTU    : 1600                    Oper Path MTU     : 1600
Far End           : 1.1.1.1                 Delivery          : GRE

Admin State       : Up                      Oper State        : Up
Acct. Pol         : None                    Collect Stats     : Disabled
Ingress Label     : 115066                  Egress Label      : 119068
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  : Preferred              Oper ControlWord  : True
Last Status Change : 02/05/2007 16:39:22    Signaling         : TLDP
Last Mgmt Change   : 02/05/2007 16:39:22
Endpoint          : N/A                     Precedence        : 4
Flags             : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : pwe3ControlWord
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
E. Fwd. Pkts.     : 0                       E. Fwd. Octets    : 0

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALA-Dut-B>config>service>epipe#
```

The following is an example when one side (PE) has the control word enabled (the pipe will be down):

This is the side with control word disabled:

```
*A:ALA-Dut-B>config>service>epipe# show service id 2100 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 1:2001  -(1.1.1.1)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id            : 1:2001                  Type              : Spoke
VC Type           : Ether                   VC Tag            : n/a
Admin Path MTU    : 1600                    Oper Path MTU     : 1600
Far End           : 1.1.1.1                 Delivery          : GRE

Admin State       : Up                      Oper State        : Down
Acct. Pol         : None                    Collect Stats     : Disabled
Ingress Label     : 115066                  Egress Label      : 119068
Ing mac Fltr      : n/a                     Egr mac Fltr      : n/a
Ing ip Fltr       : n/a                     Egr ip Fltr       : n/a
Ing ipv6 Fltr     : n/a                     Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred          Oper ControlWord  : False
Last Status Change : 02/05/2007 16:47:54    Signaling         : TLDP
Last Mgmt Change   : 02/05/2007 16:47:54
Endpoint          : N/A                     Precedence        : 4
Flags             : ReleasedIngVCLabel
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : pwe3ControlWord
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
E. Fwd. Pkts.     : 0                       E. Fwd. Octets    : 0

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALA-Dut-B>config>service>epipe#
```

This is the side with control word enabled:

```
*A:ALA-Dut-B# show service id 2100 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 1:12000  -(3.3.3.3)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id            : 1:12000                 Type              : Spoke
VC Type           : Ether                   VC Tag            : n/a
Admin Path MTU    : 1600                    Oper Path MTU     : 1600
```

```
Far End            : 3.3.3.3                 Delivery          : GRE

Admin State        : Up                      Oper State        : Down
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : 119066                  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  : Preferred               Oper ControlWord  : True
Last Status Change : 02/04/2007 22:52:43     Signaling         : TLDP
Last Mgmt Change   : 02/04/2007 02:06:08
Endpoint           : N/A                     Precedence        : 4
Flags              : NoEgrVCLabel
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
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
E. Fwd. Pkts.      : 0                       E. Fwd. Octets    : 0

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS

-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALA-Dut-B#
```

The following is an example when both sides have control word disabled:

```
*A:ALA-Dut-B>config>service>epipe# show service id 2100 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 1:2001  -(1.1.1.1)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 1:2001                  Type              : Spoke
VC Type            : Ether                   VC Tag            : n/a
Admin Path MTU     : 1600                    Oper Path MTU     : 1600
Far End            : 1.1.1.1                 Delivery          : GRE

Admin State        : Up                      Oper State        : Up
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : 115066                  Egress Label      : 119068
Ing mac Fltr       : n/a                     Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                     Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                     Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred           Oper ControlWord  : False
Last Status Change : 02/05/2007 16:49:05     Signaling         : TLDP
Last Mgmt Change   : 02/05/2007 16:47:54
```

```
Endpoint           : N/A                      Precedence        : 4
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : None
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
E. Fwd. Pkts.      : 0                        E. Fwd. Octets    : 0

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS

-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALA-Dut-B>config>service>epipe#
```

Clear Commands

id

Syntax **id** *service-id*

Context clear>service
clear>service>statistics

Description This command clears commands for a specific service.

Parameters *service-id* — The ID that uniquely identifies a service.

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 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* — Specifies the physical port identifier portion of the SAP definition.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280

ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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

Values 1 — 4294967295

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

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

Values 1 — 4294967295

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.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

all — Clears all SAP queue statistics and STP statistics.

counters — Clears all queue statistics associated with the SAP.

stp — Clears all STP statistics associated with the SAP.

sdp

Syntax **sdp** *sdp-id* **keep-alive**

Context clear>service>statistics

Description 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 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.

Virtual Private LAN Service

In This Chapter

This chapter provides information about Virtual Private LAN Service (VPLS), process overview, and implementation notes.

Topics in this chapter include:

VPLS Service Overview

Virtual Private LAN Service (VPLS) as described in Internet Draft *draft-ietf-ppvpn-vpls-ldp-01.txt*, is a class of virtual private network service that allows the connection of multiple sites in a single bridged domain over a provider-managed IP/MPLS network. The customer sites in a VPLS instance appear to be on the same LAN, regardless of their location. VPLS uses an Ethernet interface on the customer-facing (access) side which simplifies the LAN/WAN boundary and allows for rapid and flexible service provisioning.

VPLS offers a balance between point-to-point Frame Relay service and outsourced routed services (VPRN). VPLS enables each customer to maintain control of their own routing strategies. All customer routers in the VPLS service are part of the same subnet (LAN) which simplifies the IP addressing plan, especially when compared to a mesh constructed from many separate point-to-point connections. The VPLS service management is simplified since the service is not aware of nor participates in the IP addressing and routing.

A VPLS service provides connectivity between two or more SAPs on one (which is considered a local service) or more (which is considered a distributed service) 7750 SR routers. The connection appears to be a bridged domain to the customer sites so protocols, including routing protocols, can traverse the VPLS service.

Other VPLS advantages include:

- VPLS is a transparent, protocol-independent service.
- There is no Layer 2 protocol conversion between LAN and WAN technologies.
- There is no need to design, manage, configure, and maintain separate WAN access equipment, thus, eliminating the need to train personnel on WAN technologies such as Frame Relay.

VPLS over MPLS

The VPLS architecture proposed in *draft-ietf-ppvpn-vpls-ldp-0x.txt* specifies the use of provider equipment (PE) that is capable of learning, bridging, and replication on a per-VPLS basis. The PE routers that participate in the service are connected using MPLS Label Switched Path (LSP) tunnels in a full-mesh composed of mesh SDPs or based on an LSP hierarchy (Hierarchical VPLS (H-VPLS)) composed of mesh SDPs and spoke SDPs.

Multiple VPLS services can be offered over the same set of LSP tunnels. Signaling specified in *draft-martini-l2circuit-encap-mpls-0x.txt* is used to negotiate a set of ingress and egress VC labels on a per-service basis. The VC labels are used by the PE routers for de-multiplexing traffic arriving from different VPLS services over the same set of LSP tunnels.

VPLS is provided over MPLS by:

- Connecting bridging-capable provider edge routers with a full mesh of MPLS LSP (label switched path) tunnels.
- Negotiating per-service VC labels using *draft-Martini* encapsulation.
- Replicating unknown and broadcast traffic in a service domain.
- Enabling MAC learning over tunnel and access ports (see VPLS MAC Learning and Packet Forwarding).
- Using a separate forwarding information base (FIB) per VPLS service.

VPLS MAC Learning and Packet Forwarding

The 7750 SR edge devices perform the packet replication required for broadcast and multicast traffic across the bridged domain. MAC address learning is performed by the 7750 SR to reduce the amount of unknown destination MAC address flooding.

7750 SR routers learn the source MAC addresses of the traffic arriving on their access and network ports. Each 7750 SR maintains a Forwarding Information Base (FIB) for each VPLS service instance and learned MAC addresses are populated in the FIB table of the service. All traffic is switched based on MAC addresses and forwarded between all participating 7750 SR using the LSP tunnels. Unknown destination packets (i.e., the destination MAC address has not been learned) are forwarded on all LSPs to the participating 7750 SR for that service until the target station responds and the MAC address is learned by the 7750 SR associated with that service.

MAC Learning Protection

In a Layer 2 environment, subscribers connected to SAPs A, B, C can create a denial of service attack by sending packets sourcing the gateway MAC address. This will move the learned gateway MAC from the uplink SDP/SAP to the subscriber's SAP causing all communication to the gateway to be disrupted. If local content is attached to the same VPLS (D), a similar attack can be launched against it. Communication between subscribers must also be disallowed but split-horizon will not be sufficient in the topology depicted in Figure 33.

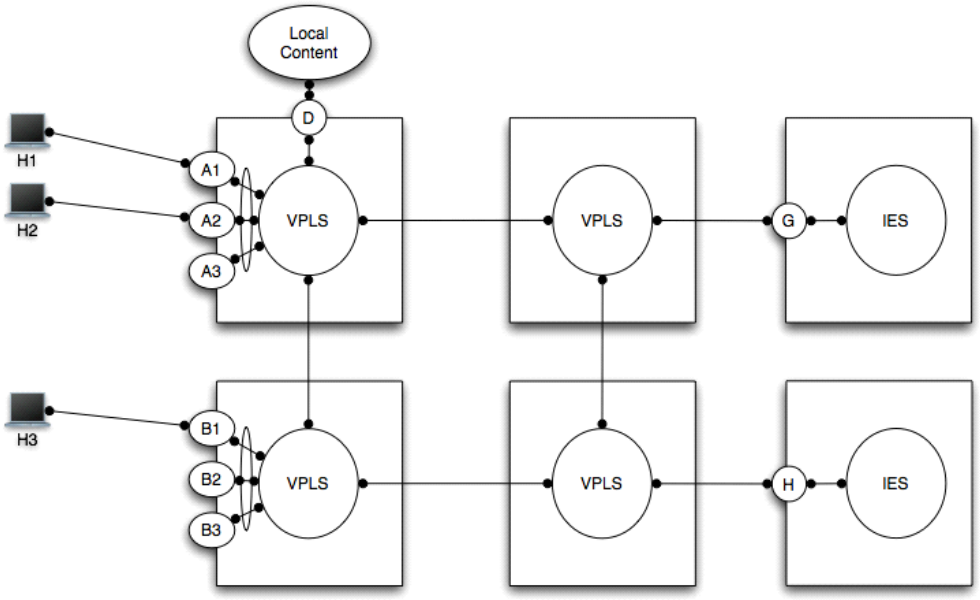

Figure 1: MAC Learning Protection

The 7750 SR-Series enables MAC learning protection capability for SAPs and SDPs. With this mechanism, forwarding and learning rules apply to the non-protected SAPs. Assume hosts H1, H2 and H3 (Figure 33) are non-protected while IES interfaces G and H are protected. When a frame arrives at a protected SAP/SDP the MAC is learned as usual. When a frame arrives from a non-protected SAP or SDP the frame must be dropped if the source MAC address is protected and the MAC address is not relearned. The system allows only packets with a protected MAC destination address.

The system may be configured the following ways:

- Static — The addresses of all protected MACs are configured. Only the IP address can be included and use a dynamic mechanism to resolve the MAC address (cpe-ping). All protected MACs in all VPLS instances in the network must be configured.
- Dynamic — The edge SAPs to protect all MAC addresses learned through them are configured. Every MAC address learned on a protection enabled SAP will be protected. The origin VPLS instance will signal the protection of the MAC to other instances in the VPN. The protection of the MAC address must be learned in context of the announcing VPN member and the MAC must be protected at that SDP. The MAC address can be relearned on a different SDP with this mechanism.
- Both — Static and dynamic.

In order to eliminate the ability of a subscriber to cause a DOS attack, the node restricts the learning of protected MAC addresses based on a statically defined list. In addition the destination MAC address is checked against the protected MAC list to verify that a packet entering a restricted SAP has a protected MAC as a destination.

VPLS Packet Walkthrough

This section provides an example of VPLS processing of a customer packet sent across the network (Figure 34) from site-A, which is connected to PE-Router-A, to site-B, which is connected to PE-Router-C (Figure 35).

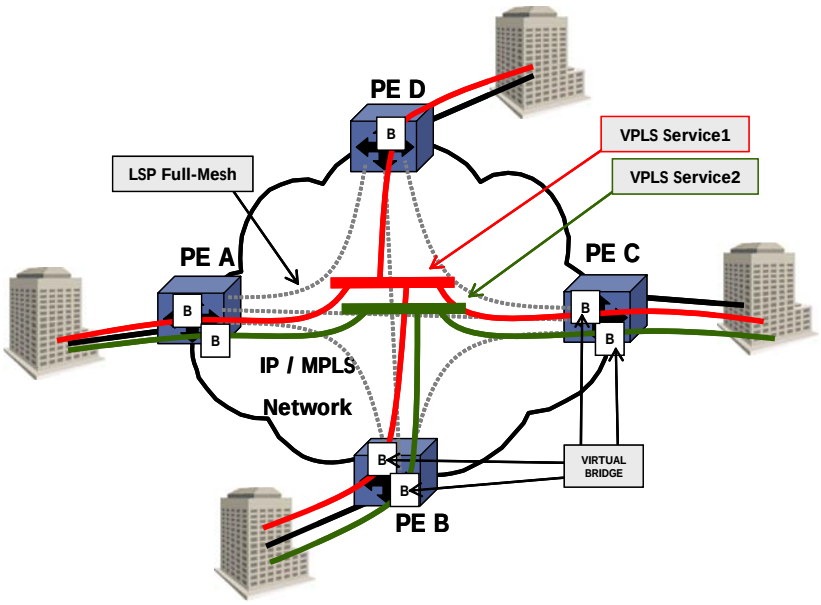

Figure 2: VPLS Service Architecture

1. PE-Router-A (Figure 35)
   a. Service packets arriving at PE-Router-A are associated with a VPLS service instance based on the combination of the physical port and the IEEE 802.1Q tag (VLAN-ID) in the packet.

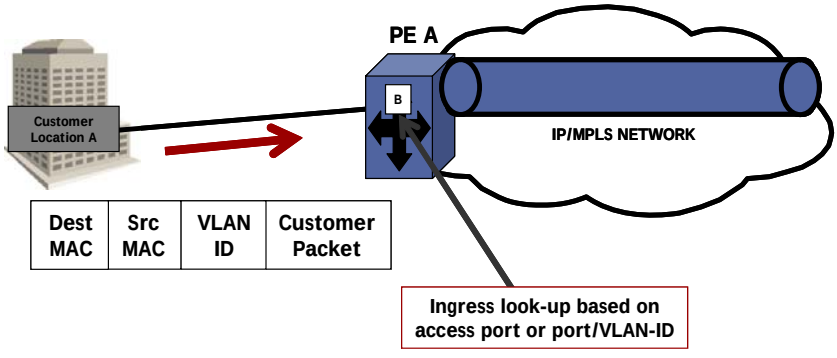

Figure 3: Access Port Ingress Packet Format and Lookup

   b. PE-Router-A learns the source MAC address in the packet and creates an entry in the FIB

table that associates the MAC address to the service access point (SAP) on which it was received.

c. The destination MAC address in the packet is looked up in the FIB table for the VPLS instance. There are two possibilities: either the destination MAC address has already been learned (known MAC address) or the destination MAC address is not yet learned (unknown MAC address).

For a Known MAC Address (Figure 36):

d. If the destination MAC address has already been learned by PE-Router-A, an existing entry in the FIB table identifies the far-end PE-Router and the service VC-label (inner label) to be used before sending the packet to far-end PE-Router-C.

e. PE-Router-A chooses a transport LSP to send the customer packets to PE-Router-C. The customer packet is sent on this LSP once the IEEE 802.1Q tag is stripped and the service VC-label (inner label) and the transport label (outer label) are added to the packet.

For an Unknown MAC Address (Figure 36):

f. If the destination MAC address has not been learned, PE-Router-A will flood the packet to both PE-Router-B and PE-Router-C that are participating in the service by using the VC-labels that each PE-Router previously signaled for the VPLS instance. Note that the packet is not sent to PE-Router-D since this VPLS service does not exist on that PE-Router.

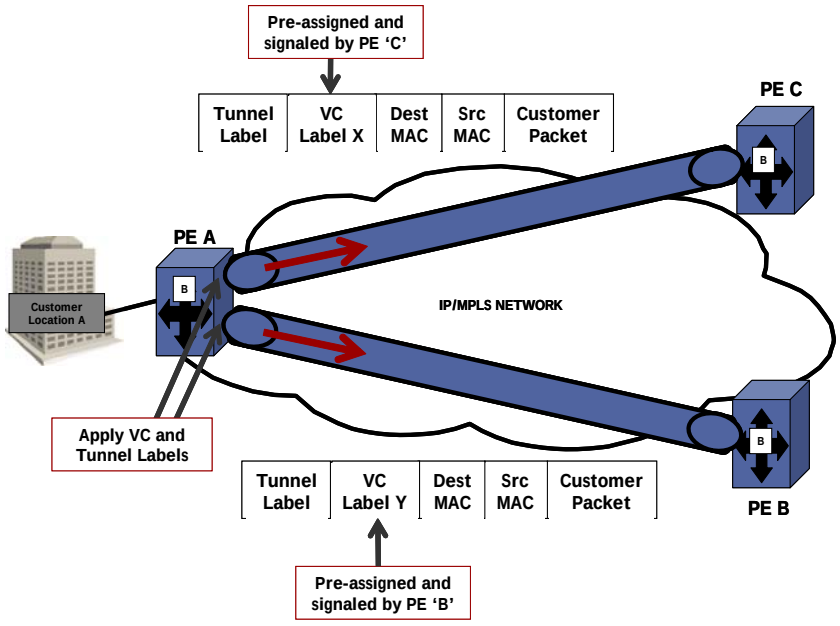

Figure 4: Network Port Egress Packet Format and Flooding

2. Core Router Switching

   a. All the core routers ('P' routers in IETF nomenclature) between PE-Router-A and PE-Router-B and PE-Router-C are Label Switch Routers (LSRs) that switch the packet based on the transport (outer) label of the packet until the packet arrives at far-end PE-Router. All core routers are unaware that this traffic is associated with a VPLS service.

3. PE-Router-C

   a. PE-Router-C strips the transport label of the received packet to reveal the inner VC-label. The VC-label identifies the VPLS service instance to which the packet belongs.

   b. PE-Router-C learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address to PE-Router-A and the VC-label that PE-Router-A signaled it for the VPLS service.

   c. The destination MAC address in the packet is looked up in the FIB table for the VPLS instance. Again, there are two possibilities: either the destination MAC address has already been learned (known MAC address) or the destination MAC address has not been learned on the access side of PE-Router-C (unknown MAC address).

      Known MAC address (Figure 37)

   d. If the destination MAC address has been learned by PE-Router-C, an existing entry in the FIB table identifies the local access port and the IEEE 802.1Q tag to be added before sending the packet to customer Location-C. The egress Q tag may be different than the ingress Q tag.

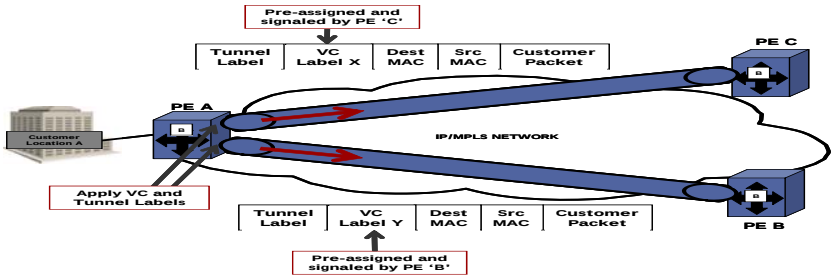

Figure 5: Access Port Egress Packet Format and Lookup

VPLS Features

This section features:

VPLS Enhancements

Alcatel-Lucent's VPLS implementation includes several enhancements beyond basic VPN connectivity. The following VPLS features can be configured individually for each VPLS service instance:

- Extensive MAC and IP filter support (up to Layer 4). Filters can be applied on a per SAP basis.
- Forwarding Information Base (FIB) management features on a per service level including:
  - → Configurable FIB size limit
  - → FIB size alarms
  - → MAC learning disable
  - → Discard unknown
  - → Separate aging timers for locally and remotely learned MAC addresses.
- Ingress rate limiting for broadcast, multicast, and destination unknown flooding on a per SAP basis.
- 7750 SR OS implementation of Spanning Tree Protocol (STP) parameters on a per VPLS, per SAP and per spoke SDP basis.
- A split horizon group on a per-SAP and per-spoke SDP basis.
- DHCP snooping and anti-spoofing on a per-SAP and per-SDP basis.
- IGMP snooping on a per-SAP and per-SDP basis.
- Optional SAP and/or spoke SDP redundancy to protect against node failure.

Table management

The following sections describe VPLS features related to management of the Forwarding Information Base (FIB).

FIB Size

The following MAC table management features are required for each instance of a SAP or spoke SDP within a particular VPLS service instance:

- MAC FIB size limits — Allows users to specify the maximum number of MAC FIB entries that are learned locally for a SAP or a spoke SDP. If the configured limit is reached, then no new addresses will be learned from the SAP until at least one FIB entry is aged out or cleared.
  - → When the limit is reached on a SAP, packets with unknown source MAC addresses are still forwarded (this default behavior can be changed via configuration). By default, if the destination MAC address is known, it is forwarded based on the FIB, and if the destination MAC address is unknown, it will be flooded. Alternatively, if "discard unknown" is enabled at the VPLS service level, any packets from unknown source MAC addresses are discarded at the SAP.
  - → The log event "SAP MAC Limit Reached" is generated when the limit is reached. When the condition is cleared, the log event "SAP MAC Limit Reached Condition Cleared" is generated.
  - → Disable learning allows users to disable the dynamic learning function on a SAP or a spoke SDP of a VPLS service instance.
  - → Disable aging allows users to turn off aging for learned MAC addresses on a SAP or a spoke SDP of a VPLS service instance.

FIB Size Alarms

The size of the VPLS FIB can be configured with a low watermark and a high watermark, expressed as a percentage of the total FIB size limit. If the actual FIB size grows above the configured high watermark percentage, an alarm is generated. If the FIB size falls below the configured low watermark percentage, the alarm is cleared by the system.

Local and Remote Aging Timers

Like a Layer 2 switch, learned MACs within a VPLS instance can be aged out if no packets are sourced from the MAC address for a specified period of time (the aging time). In each VPLS service instance, there are independent aging timers for locally learned MAC and remotely learned MAC entries in the forwarding database (FIB). A local MAC address is a MAC address associated with a SAP because it ingressed on a SAP. A remote MAC address is a MAC address received via an SDP from another 7750 SR router for the VPLS instance. The local-age timer for the VPLS instance specifies the aging time for locally learned MAC addresses, and the remote-age timer specifies the aging time for remotely learned MAC addresses.

In general, the remote-age timer is set to a longer period than the local-age timer to reduce the amount of flooding required for destination unknown MAC addresses.

Disable MAC Aging

The MAC aging timers can be disabled which will prevent any learned MAC entries from being aged out of the FIB. When aging is disabled, it is still possible to manually delete or flush learned MAC entries. Aging can be disabled for learned MAC addresses on a SAP or a spoke SDP of a VPLS service instance.

Disable MAC Learning

When MAC learning is disabled, new source MAC addresses are not entered in the VPLS FIB, whether the MAC address is local or remote. MAC learning can be disabled for individual SAPs or spoke SDPs.

Unknown MAC Discard

Unknown MAC discard is a feature which discards all packets ingressing the service where the destination MAC address is not in the FIB. The normal behavior is to flood these packets to all end points in the service.

Unknown MAC discard can be used with the disable MAC learning and disable MAC aging options to create a fixed set of MAC addresses allowed to ingress and traverse the service.

VPLS and Rate Limiting

Traffic that is normally flooded throughout the VPLS can be rate limited on SAP ingress through the use of service ingress QoS policies. In a service ingress QoS policy, individual queues can be defined per forwarding class to provide shaping of broadcast traffic, MAC multicast traffic and unknown destination MAC traffic.

MAC Move

The MAC move feature is useful to protect against undetected loops in a VPLS topology as well as the presence of duplicate MACs in a VPLS service.

If two clients in the VPLS have the same MAC address, the VPLS will experience a high re-learn rate for the MAC. When MAC move is enabled, the SR-Series will shut down the SAP or spoke SDP and create an alarm event when the threshold is exceeded.

Split Horizon SAP Groups and Split Horizon Spoke SDP Groups

Within the context of VPLS services, a loop-free topology within a fully meshed VPLS core is achieved by applying a split-horizon forwarding concept that packets received from a mesh SDP are never forwarded to other mesh SDPs within the same service. The advantage of this approach is that no protocol is required to detect loops within the VPLS core network.

In applications such as DSL aggregation, it is useful to extend this split-horizon concept also to groups of SAPs and/or spoke SDPs. This extension is referred to as a split horizon SAP group or residential bridging.

Traffic arriving on a SAP or a spoke SDP within a split horizon group will not be copied to other SAPs and spoke SDPs in the same split horizon group (but will be copied to SAPs / spoke SDPs in other split horizon groups if these exist within the same VPLS).

VPLS and Spanning Tree Protocol

Alcatel-Lucent's VPLS service emulates a bridged or switched Ethernet Layer 2 network. Equipment connected to SAPs forward Ethernet packets into the VPLS service. The 7750 SR participating in the service learns where the customer MAC addresses reside, on ingress SAPs or ingress SDPs.

Unknown destinations, broadcasts, and multicasts are flooded to all other SAPs in the service. If SAPs are connected together, either through misconfiguration or for redundancy purposes, loops can form and flooded packets can keep flowing through the network. Alcatel-Lucent's implementation of the Spanning Tree Protocol (STP) is designed to remove these loops from the VPLS topology. This is done by putting one or several SAPs and/or Spoke SDPs in the discarding state.

Alcatel-Lucent's implementation of the Spanning Tree Protocol (STP) incorporates some modifications to make the operational characteristics of VPLS more effective.

The STP instance parameters allow the balancing between resiliency and speed of convergence extremes. Modifying particular parameters can affect the behavior. For information on command usage, descriptions, and CLI syntax, refer to Configuring a VPLS Service with CLI on page 361.

Spanning Tree Operating Modes

Per VPLS instance, a preferred STP variant can be configured. The STP variants supported on the 7750 SR are:

- `rstp` — Rapid Spanning Tree Protocol (RSTP) compliant with IEEE 802.1D-2004 - default mode
- `dot1w` — Compliant with IEEE 802.1w
- `comp-dot1w` — Operation as in RSTP but backwards compatible with IEEE 802.1w (this mode was introduced for interoperability with some MTU types)
- `mstp` — Compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q-REV/D5.0-09/2005. This mode of operation is only supported in an mVPLS.

While the 7750 SR initially uses the mode configured for the VPLS, it will dynamically fall back (on a per-SAP basis) to STP (IEEE 802.1D-1998) based on the detection of a BPDU of a different format. A trap or log entry is generated for every change in spanning tree variant.

Some older 802.1W compliant RSTP implementations may have problems with some of the features added in the 802.1D-2004 standard. Interworking with these older systems is improved with the comp-dot1w mode. The differences between the RSTP mode and the comp-dot1w mode are:

- The RSTP mode implements the improved convergence over shared media feature, for example, RSTP will transition from discarding to forwarding in 4 seconds when operating over shared media. The comp-dot1w mode does not implement this 802.1D-2004 improvement and transitions conform to 802.1w in 30 seconds (both modes implement fast convergence over point-to-point links).
- In the RSTP mode, the transmitted BPDUs contain the port's designated priority vector (DPV) (conforms to 802.1D-2004). Older implementations may be confused by the DPV in a BPDU and may fail to recognize an agreement BPDU correctly. This would result in a slow transition to a forwarding state (30 seconds). For this reason, in the comp-dot1w mode, these BPDUs contain the port's port priority vector (conforms to 802.1w).

The 7750 SR supports two BDPU encapsulation formats, and can dynamically switch between these (again on a per-SAP basis):

- IEEE 802.1D STP
- Cisco PVST

Multiple Spanning Tree

The Multiple Spanning Tree Protocol (MSTP) extends the concept of the IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) by allowing grouping and associating VLANs to Multiple Spanning Tree Instances (MSTI). Each MSTI can have its own topology, which provides architecture enabling load balancing by providing multiple forwarding paths. At the same time, the number of STP instances running in the network is significantly reduced as compared to Per VLAN STP (PVST) mode of operation. Network fault tolerance is also improved because a failure in one instance (forwarding path) does not affect other instances.

The 7x50 implementation of Management VPLS (mVPLS) is used to group different VPLS instances under single RSTP instance. Introducing MSTP into the mVPLS allows interoperating with traditional Layer 2 switches in access network and provides an effective solution for dual homing of many business Layer 2 VPNs into a provider network.

Redundancy Access to VPLS

The GigE MAN portion of the network is implemented with traditional switches. Using MSTP running on individual switches facilitates redundancy in this part of the network. In order to provide dual homing of all VPLS services accessing from this part of the network, the VPLS PEs must participate in MSTP.

This can be achieved by configuring mVPLS on VPLS-PEs (only PEs directly connected to GigE MAN network) and then assign different managed-vlan ranges to different MSTP instances. Typically, the mVPLS would have SAPs with null encapsulations (to receive, send, and transmit MSTP BPDUs) and a mesh SDP to interconnect a pair of VPLS PEs.

MSTP General Principles

MSTP represents modification of RSTP which allows the grouping of different VLANs into multiple MSTIs. To enable different devices to participate in MSTIs, they must be consistently configured. A collection of interconnected devices that have the same MST configuration (region-name, revision and VLAN-to-instance assignment) comprises an MST region.

There is no limit to the number of regions in the network, but every region can support a maximum of 16 MSTIs. Instance 0 is a special instance for a region, known as the Internal Spanning Tree (IST) instance. All other instances are numbered from 1 to 4094. IST is the only spanning-tree instance that sends and receives BPDUs (typically BPDUs are untagged). All other spanning-tree instance information is included in MSTP records (M-records), which are encapsulated within MSTP BPDUs. This means that single BPDU carries information for multiple MSTI which reduces overhead of the protocol.

Any given MSTI is local to an MSTP region and completely independent from an MSTI in other MST regions. Two redundantly connected MST regions will use only a single path for all traffic flows (no load balancing between MST regions or between MST and SST region).

Traditional L2 switches running MSTP protocol assign all VLANs to the IST instance per default. The operator may then “re-assign” individual VLANs to a given MSTI by configuring per VLAN assignment. This means that a 7x50 PE can be considered as the part of the same MST region only if the VLAN assignment to IST and MSTIs is identical to the one of L2 switches in access network.

MSTP in the 7x50 Platform

The 7x50 platform uses a concept of mVPLS to group different SAPs under a single STP instance. The VLAN range covering SAPs to be managed by a given mVPLS is declared under a specific mVPLS SAP definition. MSTP mode-of-operation is only supported in an mVPLS.

When running MSTP, by default, all VLANs are mapped to the CIST. On the VPLS level VLANs can be assigned to specific MSTIs. When running RSTP, the operator must explicitly indicate, per SAP, which VLANs are managed by that SAP.

Enhancements to the Spanning Tree Protocol

To interconnect 7750 SR routers (PE devices) across the backbone, service tunnels (SDPs) are used. These service tunnels are shared among multiple VPLS instances. Alcatel-Lucent's implementation of the Spanning Tree Protocol (STP) incorporates some enhancements to make the operational characteristics of VPLS more effective. The implementation of STP on the router is modified in order to guarantee that service tunnels will not be blocked in any circumstance without imposing artificial restrictions on the placement of the root bridge within the network. The modifications introduced are fully compliant with the 802.1D-2004 STP specification.

When running MSTP, spoke SDPs cannot be configured. Also, ensure that all bridges connected by mesh SDPs are in the same region. If not, the mesh will be prevented from becoming active (trap is generated).

In order to achieve this, all mesh SDPs are dynamically configured as either root ports or designated ports. The PE devices participating in each VPLS mesh determine (using the root path cost learned as part of the normal protocol exchange) which of the 7750 SR devices is closest to the root of the network. This PE device is internally designated as the primary bridge for the VPLS mesh. As a result of this, all network ports on the primary bridges are assigned the designated port role and therefore remain in the forwarding state.

The second part of the solution is to ensure that the remaining PE devices participating in the STP instance see the SDP ports as a lower cost path to the root rather than a path that is external to the mesh. Internal to the PE nodes participating in the mesh, the SDPs are treated as zero cost paths towards the primary bridge. As a consequence, the path through the mesh will be seen as lower cost than any alternative and the PE node will designate the network port as the root port. This approach ensures that network ports will always remain in forwarding state.

In combination, these two enhancements guarantee that network ports will never be blocked and yet maintain interoperability with bridges external to the mesh which are running STP instances.

L2PT termination

L2PT is used to transparently transport protocol data units (PDUs) of L2 protocols such as STP, CDP and VTP. This allows running these protocols between customer CPEs without involving service provider infrastructure.

The 7x50 allows transparent tunneling of PDUs across the VPLS core. However, in some network designs VPLS PE is connected to CPEs through legacy Layer 2 network, rather than having direct connections. In such environments termination of tunnels through such infrastructure is required.

L2PT tunnels protocol PDUs by overwriting MAC destination addresses at the ingress of the tunnel to a proprietary MAC address such as 01-00-0c-cd-cd-d0. At the egress of the tunnel, this MAC address is then overwritten back to MAC address of the respective Layer 2 protocol.

7x50 supports L2PT termination for STP BPDUs only. More specifically:

- At ingress of every SAP/spoke SDP which is configured as L2PT termination, all PDUs with a MAC destination address, 01-00-0c-cd-cd-d0 will be intercepted and their MAC destination address will be overwritten to MAC destination address used for the corresponding protocol (PVST, STP, RSTP). The type of the STP protocol can be derived from LLC and SNAP encapsulation.
- In egress direction, all STP PDUs received on all VPLS ports will be intercepted and LT2P encapsulation will be performed for SAP/spoke SDPs configured as L2PT termination points. Because of the implementation reasons, PDU interception and re-direction to CPM can be performed only at ingress. Therefore, to comply with the above requirement, as soon as at least 1 port of a given VPLS service is configured as L2PT termination port, redirection of PDUs to CPM will be set on all other ports (SAPs, spoke SDPs and mesh SDPs) of the VPLS service.

L2PT termination can be enabled only if STP is disabled in a context of the given VPLS service.

BPDU Translation

VPLS networks are typically used to interconnect different customer sites using different access technologies such as Ethernet and bridged-encapsulated ATM PVCs. Typically, different Layer 2 devices can support different types of STP and even if they are from the same vendor. In some cases it is therefore necessary to provide BPDU translation in order to provide interoperable e2e a solution.

To address these network designs, BPDU format translation is supported on 7x50 devices. If enabled on a given SAP or spoke SDP, the system will intercept all BPDUs destined to that interface and perform required format translation such as STP-to-PVST or vice versa.

Similarly, BPDU interception and redirection to the CPM is performed only at ingress meaning that as soon as at least 1 port within a given VPLS service has BPDU translation enabled, all BPDUs received on any of this VPLS ports will be redirected to the CPM.

BPDU translation involves all encapsulation actions that the data path would do for a given outgoing port (adding VLAN tags, depending on outer SAP, SDP encapsulation type) and adding or removing all required VLAN information in a BPDU payload.

This feature can be enabled on a SAP only if STP is disabled in the context of the given VPLS service.

Egress Multicast Groups

Efficient multicast replication is a method of increasing egress replication performance by combining multiple destinations into a single egress forwarding pass. In standard egress VPLS multicast forwarding, the complete egress forwarding plane is used per destination to provide ACL, mirroring, QoS and accounting for each path with associated receivers. In order to apply the complete set of available egress VPLS features, the egress forwarding plane must loop-back copies of the original packet so that each flooding destination may be processed. While each distributed egress forwarding plane only replicates to the destinations currently reached through its ports, this loop-back and replicate function can be resource intensive. When egress forwarding plane congestion conditions exist, unicast discards may be indiscriminate relative to forwarding priority. Another by-product of this approach is that the ability for the forwarding plane to fill the egress links is affected which could cause under-run conditions on each link while the forwarding plane is looping packets back to itself.

In an effort to provide highly scalable VPLS egress multicast performance for triple play type deployments, an alternative efficient multicast forwarding option is being offered. This method allows the egress forwarding plane to send a multicast packet to a set (called a chain) of destination SAPs with only a single pass through the egress forwarding plane. This minimizes the egress resources (processing and traffic management) used for the set of destinations and allows proper handling of congestion conditions and minimizes line under-run events. However, due to the batch nature of the egress processing, the chain of destinations must share many attributes. Also, egress port and ACL mirroring will be disallowed for packets handled in this manner.

Packets eligible for forwarding by SAP chaining are VPLS flooded packets (broadcast, multicast and unknown destination unicast) and IP multicast packets matching an VPLS Layer 2 (s,g) record (created through IGMP snooping).

Egress Multicast Group Provisioning

To identify SAPs in the chassis that are eligible for egress efficient multicast SAP chaining, an egress multicast group must be created. SAPs from multiple VPLS contexts may be placed in a single group to minimize the number of groups required on the system and to support multicast VPLS registration (MVR) functions.

Some of the parameters associated with the group member SAPs must be configured with identical values. The common parameters are checked as each SAP is provisioned into the group. If the SAP fails to be consistent in one or more parameters, the SAP is not allowed into the egress multicast group. Once a SAP is placed into the group, changing of a common parameter is not permitted.

Required Common SAP Parameters

Only SAPs created on Ethernet ports are allowed into an egress multicast group.

Required common parameters include:

- SAP Port Encapsulation Type on page 342
- SAP Port Dot1Q EtherType on page 342
- SAP Egress Filter on page 342

SAP Port Encapsulation Type

The access port encapsulation type defines how the system will delineate SAPs from each other on the access port. SAPs placed in the egress multicast group must be of the same type. The supported access port encapsulation types are null and dot1q. While all SAPs within the egress multicast group share the same encapsulation type, they are allowed to have different encapsulation values defined. The chained replication process will make the appropriate dot1q value substitution per destination SAP.

The normal behavior of the system is to disallow changing the port encapsulation type once one or more SAPs have been created on the SAP. This being the case, no special effort is required to ensure that a SAP will be changed from null to dot1q or dot1q to null while the SAP is a member of a egress multicast group. Deleting the SAP will automatically remove the SAP from the group.

SAP Port Dot1Q EtherType

The access port dot1q-etype parameter defines which EtherType will be expected in ingress dot1q encapsulated frames and the EtherType that will be used to encapsulate egress dot1q frames on the port. SAPs placed in the same egress multicast group must use the same EtherType when dot1q is enabled as the SAPs encapsulation type.

The normal behavior of the system is to allow dynamic changing of the access port dot1q-etype value while SAPs are currently using the port. Once a dot1q SAP on an access port is allowed into an egress multicast group, the port on which the SAP is created will not accept a change of the configured dot1q-etype value. When the port encapsulation type is set to null, the ports dot1q-etype parameter may be changed at any time.

SAP Egress Filter

Due to the chaining nature of egress efficient multicast replication, only the IP or MAC filter defined for the first SAP on each chain is used to evaluate the packet. To ensure consistent behavior for all SAPs in the egress multicast group, when an IP or MAC filter is configured on one SAP it must be configured on all. To prevent inconsistencies, each SAP must have the same egress

IP or MAC filter configured (or none at all) prior to allowing the SAP into the egress multicast group.

Attempting to change the egress filter configured on the SAP while the SAP is a member of an egress multicast group is not allowed.

If the configured common egress filter is changed on the egress multicast group, the egress filter on all member SAPs will be overwritten by the new defined filter. If the SAP is removed from the group, the previous filter definition is not restored.

SAP Egress QoS Policy

Each SAP placed in the egress multicast group may have a different QoS policy defined. When the egress forwarding plane performs the replication for each destination in a chain, the internal forwarding class associated with the packet is used to map the packet to an egress queue on the SAP.

In the case where subscriber SLA management is enabled on the SAP and the SAP queues are not available, the queues created by the non-sub-addr-traffic SLA-profile instance are used.

One caveat is that egress Dot1P markings for dot1q SAPs in the replication chain are only evaluated for the first SAP in the chain. If the first SAP defines an egress Dot1P override for the packet, all encapsulations in the chain will share the same value. If the first SAP in the chain does not override the egress Dot1P value, either the existing Dot1P value (relative to ingress) will be preserved or the value 0 (zero) will be used for all SAPs in the replication chain. The egress QoS policy Dot1P remark definitions on the other SAPs in the chain are ignored by the system.

Efficient Multicast Egress SAP Chaining

The egress IOM (Input Output Module) automatically creates the SAP chains on each egress forwarding plane (typically all ports on an MDA are part of a single forwarding plane except in the case of the 10 Gigabit IOM which has two MDAs on a single forwarding plane). The size of each chain is based on the dest-chain-limit command defined on the egress multicast group to which the SAPs in the chain belong.

A set of chains is created by the IOM for each egress flooding list managed by the IOM. While SAPs from multiple VPLS contexts are allowed into a single egress multicast group, an egress flooding list is typically based on a subset of these SAPs. For instance, the broadcast/multicast/unknown flooding list for a VPLS context is limited to the SAPs in that VPLS context. With IGMP snooping on a single VPLS context, the flooding list is per Layer 2 IGMP (s,g) record and is basically limited to the destinations where IGMP joins for the multicast stream have been intercepted. When MVR (Multicast VPLS Registration) is enabled, the (s,g) flooding list may include SAPs from various VPLS contexts based on MVR configuration.

The system maintains a unique flooding list for each forwarding plane VPLS context (see section VPLS Broadcast/Multicast/Unknown Flooding List on page 344). This list will contain all SAPs (except for residential SAPs), spoke SDP and mesh SDP bindings on the forwarding plane that belong to that VPLS context. Each list may contain a maximum of 127 SAPs. In the case where the IOM is able to create an egress multicast chain, the SAPs within the chain are represented in the flooding list by a single SAP entry (the first SAP in the chain).

The system also maintains a unique flooding list for each Layer 2 IP multicast (s,g) record created through IGMP snooping (see sections VPLS IGMP Snooping (s,g) Flooding List on page 345 and MVR IGMP Snooping (s,g) Flooding List on page 345). A flooding list created by IGMP snooping is limited to 127 SAPs, although it may contain other entries representing spoke and mesh SDP bindings. Unlike a VPLS flooding list, a residential SAP may be included in a Layer 2 IP multicast flooding list.

While the system may allow 30 SAPs in a chain, the uninterrupted replication to 30 destinations may have a negative effect on other packets waiting to be processed by the egress forwarding plane. Most notably, massive jitter may be seen on real time VoIP or other time-sensitive applications. The dest-chain-limit parameter should be tuned to allow the proper balance between multicast replication efficiency and the effect on time sensitive application performance. It is expected that the optimum performance for the egress forwarding plane will be found at around 16 SAPs per chain.

VPLS Broadcast/Multicast/Unknown Flooding List

The IOM includes all VPLS destinations in the egress VPLS Broadcast/Multicast/Unknown (BMU) flooding list that exist on a single VPLS context. Whenever a broadcast, multicast or unknown destination MAC is received in the VPLS, the BMU flooding list is used to flood the packet to all destinations. For normal flooding, care is taken at egress to ensure that the packet is not sent back to the source of the packet. Also, if the packet is associated with a split horizon group (mesh or spoke/SAP) the egress forwarding plane will prevent the packet from reaching destinations in the same split horizon context as the source SAP or SDP-binding.

The VPLS BMU flooding list may contain both egress multicast group SAPs and other SAPs or SDP bindings as destinations. The egress IOM will separate the egress multicast group SAPs from the other destinations to create one or more chains. Egress multicast group SAPs are placed into a chain completely at the discretion of the IOM and the order of SAPs in the list will be nondeterministic. When more SAPs exist on the VPLS context within the egress multicast group then are allowed in a single chain, multiple SAP chains will be created. The IOM VPLS egress BMU flooding list will then contain the first SAP in each chain plus all other VPLS destinations.

The SAPs in the same VPLS context must be in the same split horizon group to allow membership into the egress multicast group. The split horizon context is not required to be the same between VPLS contexts.

SAPs within the same VPLS context may be defined in different egress multicast groups, but SAPs in different multicast groups cannot share the same chain.

VPLS IGMP Snooping (s,g) Flooding List

When IGMP snooping is enabled on a VPLS context, a Layer 2 IP multicast record (s,g) is created for each multicast stream entering the VPLS context. Each stream should only be sent to each SAP or SDP binding where either a multicast router exists or a host exists that has requested to receive the stream (known as a receiver). To facilitate egress handling of each stream, the IOM creates a flooding list for each (s,g) record associated with the VPLS context. As with the BMU flooding list, source and split horizon squelching is enforced by the egress forwarding plane.

As with the BMU VPLS flooding list, the egress multicast group SAPs that have either static or dynamic multicast receivers for the (s,g) stream are chained into groups. The chaining is independent of other (s,g) flooding lists and the BMU flooding list on the VPLS instance. As the (s,g) flooding list membership is dynamic, the egress multicast group SAPs in chains in the list are also managed dynamically.

Since all SAPs placed into the egress multicast group for a particular VPLS context are in the same split horizon group, no special function is required for split horizon squelching.

MVR IGMP Snooping (s,g) Flooding List

When IGMP snooping on a SAP is tied to another VPLS context to facilitate cross VPLS context IP multicast forwarding, a Layer 2 IP multicast (s,g) record is maintained on the VPLS context receiving the multicast stream. This is essentially an extension to the VPLS IGMP snooped flooding described in VPLS IGMP Snooping (s,g) Flooding List on page 345. The (s,g) list is considered to be owned by the VPLS context that the multicast stream will enter. Any SAP added to the list that is outside the target VPLS context (using the from-vpls command) is handled as an alien SAP. Split horizon squelching is ignored for alien SAPs.

When chaining the egress multicast group SAPs in an MVR (s,g) list, the IOM will keep the native chained SAPs in separate chains from the alien SAPs to prevent issues with split horizon squelching.

Mirroring and Efficient Multicast Replication

As previously stated, efficient multicast replication affects the ability to perform mirroring decisions in the egress forwarding plane. In the egress forwarding plane, mirroring decisions are performed prior to the egress chain replication function. Since mirroring decisions are only evaluated for the first SAP in each chain, applying a mirroring condition to packets that egress other SAPs in the chain has no effect. Also, the IOM manages the chain membership automatically and the user has no ability to provision which SAP is first in a chain. Thus, in this release, mirroring is not allowed for SAPs within a chain.

Port Mirroring

A SAP created on an access port that is currently defined as an egress mirror source may not be defined into an egress multicast group.

A port that has a SAP defined in an egress multicast group may not be defined as an egress mirror source. If egress port mirroring is desired, then all SAPs on the port must first be removed from all egress multicast groups.

Filter Mirroring

An IP or MAC filter that is currently defined on an egress multicast group as a common required parameter may not have an entry from the list defined as a mirror source.

An IP or MAC filter that has an entry defined as a mirror source may not be defined as a common required parameter for an egress multicast group.

If IP or MAC based filter mirroring is required for packets that egress an egress multicast group SAP, the SAP must first be removed from the egress multicast group and then an IP or MAC filter that is not associated with an egress multicast group must be assigned to the SAP.

SAP Mirroring

While SAP mirroring is not allowed within an IOM chain of SAPs, it is possible to define an egress multicast group member SAP as an egress mirror source. When the IOM encounters a chained SAP as an egress mirror source, it automatically removes the SAP from its chain, allowing packets that egress the SAP to hit the mirror decision. Once the SAP is removed as an egress mirror source, the SAP will be automatically placed back into a chain by the IOM.

It should be noted that all mirroring decisions affect forwarding plane performance due to the overhead of replicating the frame to the mirror destination. This is especially true for efficient multicast replication as removing the SAP from the chain also eliminates a portion of the replication efficiency along with adding the mirror replication overhead.

OAM Commands with EMG

There are certain limitations with using the OAM commands when egress multicast group (EMG) is enabled. This is because OAM commands work by looping the OAM packet back to ingress instead of sending them out of the SAP. Hence, if EMG is enabled, these OAM packets will be looped back once per chain and hence, will only be processed for the first SAP on each chain. Particularly, mac-ping, mac-trace and mfib-ping commands will only list the first SAP in each chain.

IOM Chain Management

As previously stated, the IOM automatically creates the chain lists from the available egress multicast group SAPs. The IOM will create chains from the available SAPs based on the following rules:

1. SAPs from different egress multicast groups must be in different chains (a chain can only contain SAPs from the same group)
2. Alien and native SAPs must be in different chains
3. A specific chain cannot be longer than the defined dest-chain-limit parameter for the egress multicast group to which the SAPs belong

Given the following conditions for an IOM creating a multicast forwarding list (List 1) for a Layer 2 IP multicast (s,g) native to VPLS instance 100:

- Egress multicast group A
  - → Destination chain length = 16
  - → 30 member SAPs on VPLS 100 joined (s,g) (native to VPLS 100)
  - → 41 member SAPs on other VPLS instances joined (s,g) (alien to VPLS 100)
- Egress multicast group B
  - → Destination chain length = 8
  - → 17 member SAPs on VPLS 100 joined (s,g) (native to VPLS 100)
- Egress multicast group C
  - → Destination chain length = 12
  - → 23 member SAPs on other VPLS instances joined (s,g) (alien to VPLS 100)

The system will build the SAP chains for List 1 according to Table 12.

Table 1: SAP Chain Creation

Egress Forwarding List 1 SAP Chains					
Egress Multicast Group A Destination Chain Length 16		**Egress Multicast Group B Destination Chain Length 8**		**Egress Multicast Group C Destination Chain Length 12**	
Native Chains	**Alien Chains**	**Native Chains**	**Alien Chains**	**Native Chains**	**Alien Chains**
16	16	8			12
14	16	8			11
	9	1			

Adding a SAP to a Chain

A SAP must meet all the following conditions to be chained in a VPLS BMU flooding list:

1. The SAP is successfully defined as an egress multicast group member
2. The SAP is not currently an egress mirror source

Further, a SAP must meet the following conditions to be chained in an egress IP multicast (s,g) Flooding List:

1. The SAP is participating in IGMP snooping
2. A static or dynamic join to the (s,g) record exists for the SAP or the SAP is defined as a multicast router port

Note: While an operationally down SAP is placed into replication chains, the system ignores that SAP while in the process of replication.

Based on the egress multicast group and the native or alien nature of the SAP in the list, a set of chains are selected for the SAP. The IOM will search the chains for the first empty position in an existing chain and place the SAP in that position. If an empty position is not found, the IOM will create a new chain with that SAP in the first position and add the SAP to the flooding list to represent the new chain.

Removing a SAP from a Chain

A SAP will be removed from a chain in a VPLS BMU flooding list or egress IP multicast (s,g) flooding list for any of the following conditions:

1. The SAP is deleted from the VPLS instance
2. The SAP is removed from the egress multicast group of which it was a member
3. The SAP is defined as an egress mirror source

Further, a SAP will be removed from an egress IP multicast (s,g) flooding list for the following conditions:

1. IGMP snooping removes the SAP as an (s,g) destination or the SAP is removed as a multicast router port

When the SAP is only being removed from the efficient multicast replication function, it may still need to be represented as a stand alone SAP in the flooding list. If the removed SAP is the first SAP in the list, the second SAP in the list is added to the flooding list when the first SAP is de-chained. If the removed SAP is not the first SAP, it is first de-chained and then added to the flooding list. If the removed SAP is the only SAP in the chain, the chain is removed along with removing the SAP from the flooding list.

Moving a SAP from a chain to a stand alone condition or from a stand alone condition to a chain may cause a momentary glitch in the forwarding plane for the time that the SAP is being moved. Care is taken to prevent or minimize the possibility of duplicate packets being replicated to a destination while the chains and flooding lists are being manipulated.

Chain Optimization

Chains are only dynamically managed during SAP addition and removal events. The system does not attempt to automatically optimize existing chains. It is possible that excessive SAP removal may cause multiple chains to exist with lengths less than the maximum chain length. For example, if four chains exist with eight SAPs each, it is possible that seven of the SAPs from each chain are removed. The result would be four chains of one SAP each effectively removing any benefit of egress SAP replication chaining.

While it may appear that optimization would be beneficial each time a SAP is removed, this is not the case. Rearranging the chains each time a SAP is removed may cause either packet duplication or omitting replication to a destination SAP. Also, it could be argued that if the loop back replication load is acceptable before the SAP is removed, continuing with the same loop back replication load once the SAP is removed is also acceptable. It is important to note that the overall replication load is lessened with each SAP removal from a chain.

While dynamic optimization is not supported, a manual optimization command is supported in each egress multicast group context. When executed the system will remove and add each SAP, rebuilding the replication chains.

When the dest-chain-limit is modified for an egress multicast group, the system will reorganize the replication chains that contain SAPs from that group according to the new maximum chain size.

IOM Mode B Capability

Efficient multicast replication uses an egress forwarding plane that supports chassis mode b due to the expanded memory requirements to store the replication chain information. The system does not need to be placed into mode b for efficient multicast replication to be performed. Any IOM that is capable of mode “b” operation automatically performs efficient multicast replication when a flooding list contains SAPs that are members of an egress multicast group.

VPLS Redundancy

The VPLS draft standard (draft-ietf-l2vpn-vpls-ldp, *Virtual Private LAN Services over MPLS*) includes provisions for hierarchical VPLS, using point-to-point spoke SDPs. Two applications have (so far) been identified for spoke SDPs:

- to connect to Multi-Tenant Units (MTUs) within a metro area network;
- to interconnect the VPLS nodes of two metro networks.

In both applications the spoke SDPs serve to improve the scalability of VPLS.
While node redundancy is implicit in non-hierarchical VPLS services (using a full mesh of SDPs between PEs), node redundancy for spoke SDPs needs to be provided separately.

Alcatel-Lucent routers have implemented special features for improving the resilience of Hierarchical VPLS instances, in both MTU and inter-metro applications.

Spoke SDP Redundancy for Metro Interconnection

When two or more meshed VPLS instances are interconnected by redundant spoke SDPs (as shown in Figure 38), a loop in the topology results. In order to remove such a loop from the topology, Spanning Tree Protocol (STP) can be run over the SDPs (links) which form the loop such that one of the SDPs is blocked. As running STP in each and every VPLS in this topology is not efficient, the 7750 SR includes functionality which can associate a number of VPLSes to a single STP instance running over the redundant-SDPs. Node redundancy is thus achieved by running STP in one VPLS, and applying the conclusions of this STP to the other VPLSes. The VPLS instance running STP is referred to as the 'management VPLS' or mVPLS.

In the case of a failure of the active node, STP on the management VPLS in the standby node will change the link states from disabled to active. The standby node will then broadcast a MAC flush LDP control message in each of the protected VPLS instances, so that the address of the newly active node can be re-learned by all PEs in the VPLS.

It is possible to configure two management VPLSes, where both VPLSes have different active spokes (this is achieved by changing the path-cost in STP). By associating different user VPLSes with the two management VPLSes, load balancing across the spokes can be achieved.

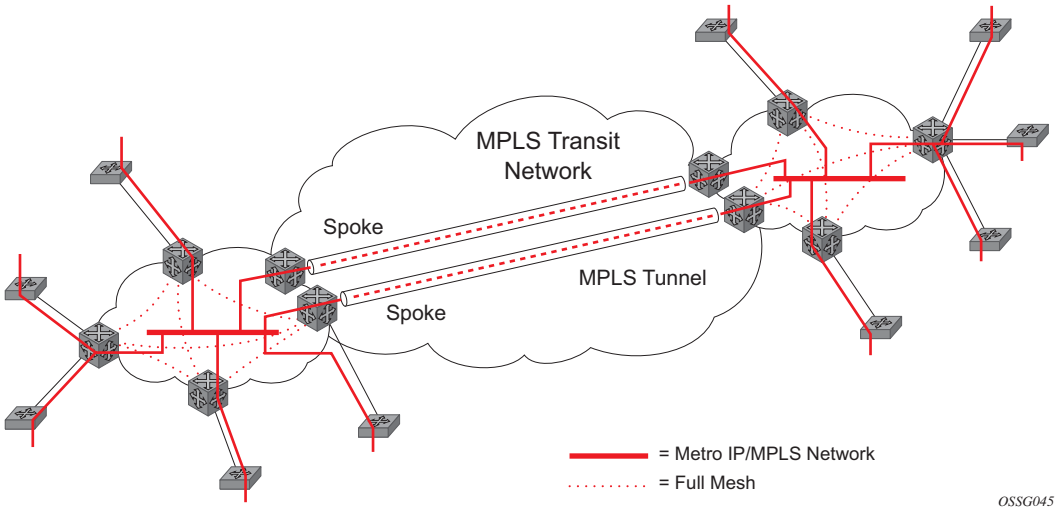

Figure 6: HVPLS with Spoke Redundancy

SAP Redundancy for MTU Protection

MAC Flush with STP

A second application of Hierarchical VPLS is in the use of Multi Tenant Units (MTU). MTUs are typically not MPLS-enabled, and thus have Ethernet links to the closest PE node (see Figure 39 below). To protect against failure of the PE node, an MTU could be dual-homed and thus have two SAPs on two PE nodes. To resolve the potential loop, STP is activated on the MTU and the two PEs.

Like in the scenario above, STP only needs to run in a single VPLS instance, and the results of the STP calculations are applied to all VPLSes on the link. Equally, the standby node will broadcast MAC flush LDP messages in the protected VPLS instances when it detects that the active node has failed.

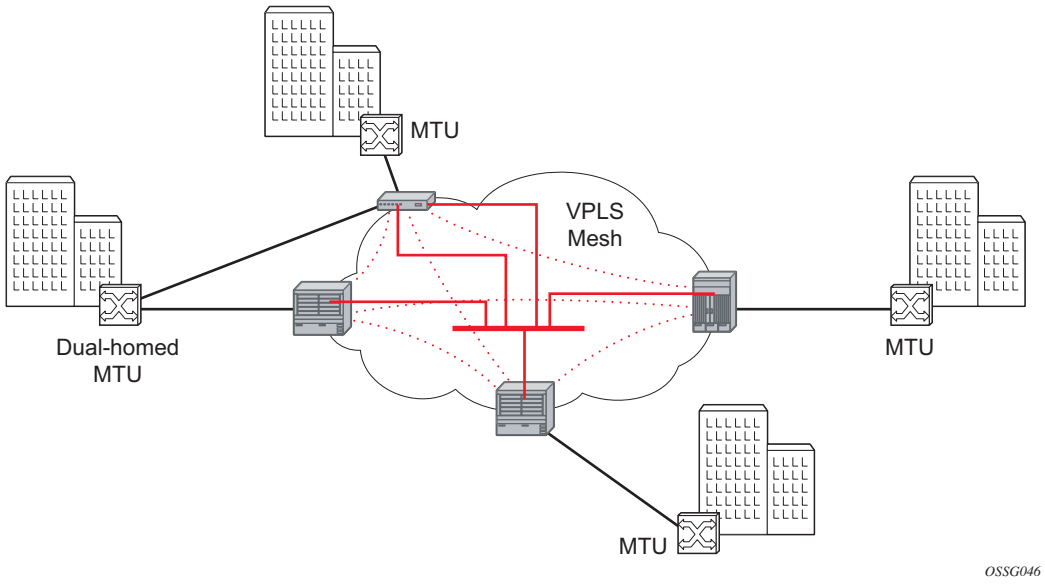

Figure 7: HVPLS with SAP Redundancy

Selective MAC flush

When using STP as described above is not appropriate, the “Selective MAC flush” feature can be used instead.

In this scenario, the 7x50 that detects a port failure will send out a flush-all-from-ME LDP message to all PEs in the VPLS. The PEs receiving this LDP message will remove all MAC entries originated by the sender from the indicated VPLS.

A drawback of this approach is that selective MAC flush itself does not signal that a backup path was found, only that the previous path is no longer available. In addition, the selective MAC Flush mechanism is effective only if the CE and PE are directly connected (no intermediate hubs or bridges) as it reacts only to a physical failure of the link. Consequently it is recommended to use the MAC flush with STP method described above where possible.

Dual Homing to a VPLS Service

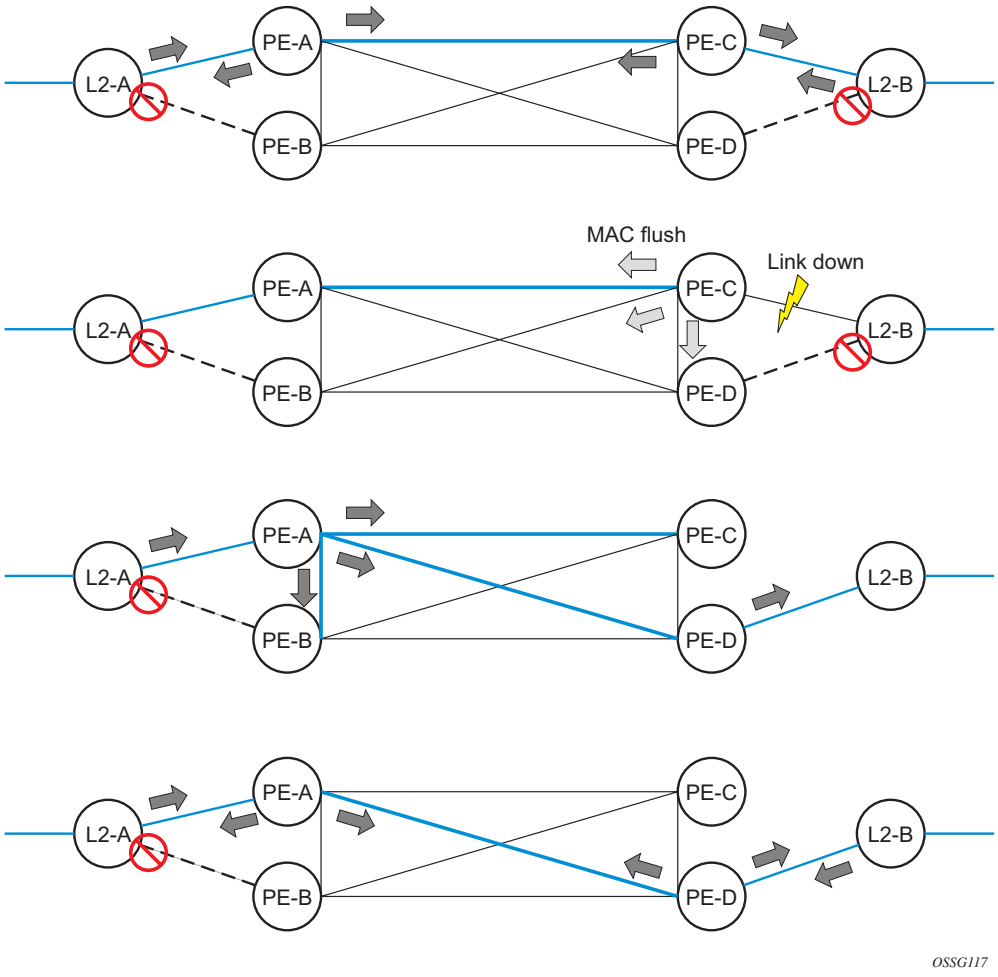

Figure 8: Dual Homed CE Connection to VPLS

Figure 40 illustrates a dual-homed connection to VPLS service (PE-A, PE-B, PE-C, PE-D) and operation in case of link failure (between PE-C and L2-B). Upon detection of a link failure PE-C will send MAC-Address-Withdraw messages, which will indicate to all LDP peers that they should flush all MAC addresses learned from PE-C. This will lead that to a broadcasting of packets addressing affected hosts and re-learning process in case alternative route exists.

Note that the message described here is different than the message described in draft-ietf-l2vpn-vpls-ldp-xx.txt, *Virtual Private LAN Services over MPLS*. The difference is in the interpretation and action performed in the receiving PE. According the draft definition, upon receipt of a MAC-withdraw message, all MAC addresses, except the ones learned from the source PE, are flushed,

This section specifies that all MAC addresses learned from the source are flushed. This message has been implemented as an LDP address message with vendor-specific type, length, value (TLV), and is called the flush-all-from-ME message.

The draft definition message is currently used in management VPLS which is using RSTP for recovering from failures in Layer 2 topologies. The mechanism described in this document represent an alternative solution.

The advantage of this approach (as compared to RSTP based methods) is that only MAC-affected addresses are flushed and not the full forwarding database. While this method does not provide a mechanism to secure alternative loop-free topology, the convergence time is dependent on the speed of the given CE device will open alternative link (L2-B switch in Figure 40) as well as on the speed PE routers will flush their FDB.

In addition, this mechanism is effective only if PE and CE are directly connected (no hub or bridge) as it reacts to physical failure of the link.

ACL Next-Hop for VPLS

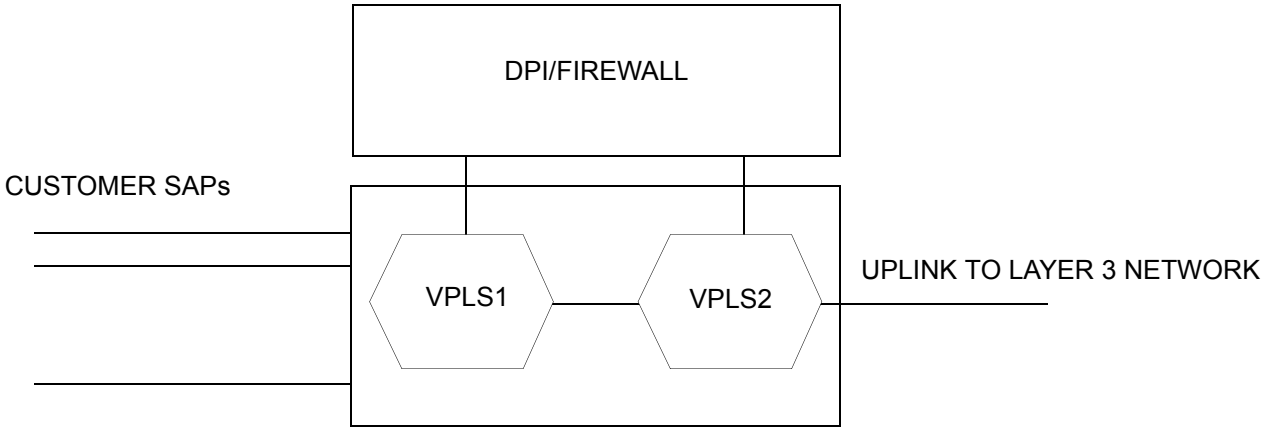

Figure 9: Application 1 Diagram

The ACL next-hop for VPLS feature enables an ACL that has a forward next-hop SAP or SDP action specified to be used in a VPLS service to direct traffic with specific match criteria to a SAP or SDP. This allows traffic destined to the same gateway to be split and forwarded differently based on the ACL.

Policy routing is a popular tool used to direct traffic in Layer 3 networks. As Layer 2 VPNs become more popular, especially in network aggregation, policy forwarding is required. Many providers are using methods such as DPI servers, transparent firewalls or Intrusion Detection/ Prevention Systems (IDS/IPS). Since these devices are bandwidth limited providers want to limit traffic forwarded through them. A mechanism is required to direct some traffic coming from a SAP to the DPI without learning and other traffic coming from the same SAP directly to the gateway uplink based learning. This feature will allow the provider to create a filter that will forward packets to a specific SAP or SDP. The packets are then forwarded to the destination SAP regardless of learned destination or lack thereof. The SAP can either terminate a Layer 2 firewall, deep packet inspection (DPI) directly or may be configured to be part of a cross connect bridge into another service. This will be useful when running the DPI remotely using VLLs. If an SDP is used the provider can terminate it in a remote VPLS or VLL service where the firewall is connected. The filter can be configured under a SAP or SDP in a VPLS service. All packets (unicast, multicast, broadcast and unknown) can be delivered to the destination SAP/SDP.

The filter may be associated SAPs/SDPs belonging to a VPLS service only if all actions in the ACL forward to SAPs/SDPs that are within the context of that VPLS. Other services (IES, VLL and VPRN) do not support this feature. An ACL that contains this feature is allowed but the system will drop any packet that matches an entry with this action.

SDP Statistics for VPLS and VLL Services

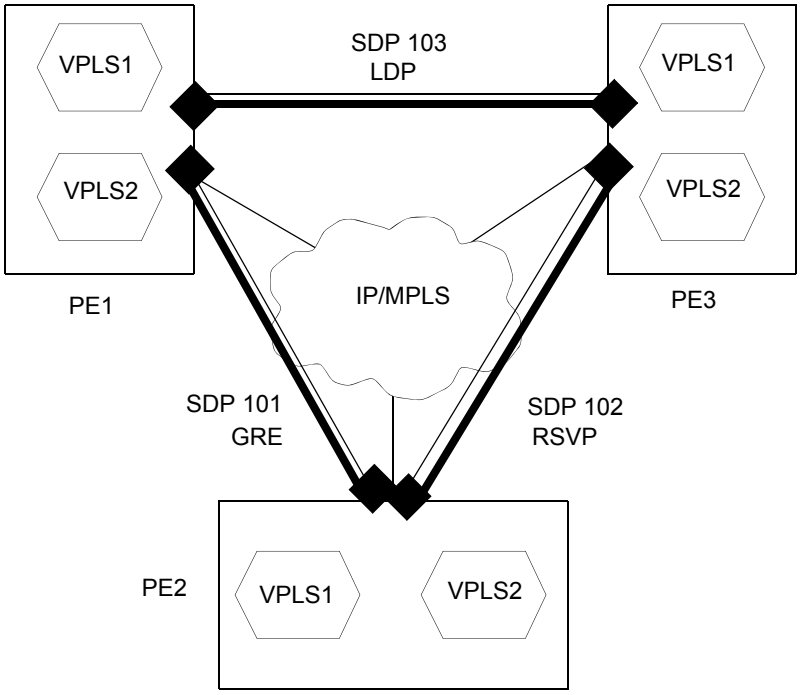

Figure 10: SDP Statistics for VPLS and VLL Services

The simple three-node network described in Figure 42 shows two MPLS SDPs and one GRE SDP defined between the nodes. These SDPs connect VPLS1 and VPLS2 instances that are defined in the three nodes. With this feature the operator will have local CLI based as well as SNMP based statistics collection for each VC used in the SDPs. This will allow for traffic management of tunnel usage by the different services and with aggregation the total tunnel usage.

SDP statistics allow providers to bill customers on a per-SDP per-byte basis. This destination-based billing model is can be used by providers with a variety of circuit types and have different costs associated with the circuits. An accounting file allows the collection of statistics in a bulk manner.

Auto SDPs and Auto SDP Bindings

When a VPLS service is created and the **radius-discovery** command is enabled, manual configuration of SDPs is not required. The far-end SDP is discovered by RADIUS and an auto-SDP will be created for each far-end PE. An auto-SDP binding will be created for the service (the SDP binding is on a per service basis). The auto-SDP can be used for other VPLS service instances between the pair of PEs. If a manual SDP binding is created for the service, it will be preferred to the auto-SDP binding to allow override.

An auto-SDP is created by a PE router automatically with information provided by RADIUS discovery. The binding, auto-SDP binding is created automatically for an auto-SDP. An auto-SDP can be GRE based or LDP based. The default is GRE.

An SNMP trap is generated to inform the NMS about the creation of an auto-SDP or auto-SDP binding. Auto-SDPs and auto-SDP bindings are not saved in the router configuration file. Auto-SDPs cannot be manually configured nor modified.

Manual SDPs and Manual SDP Bindings

"Manual" SDPs can be used with RADIUS discovery. Manual SDPs take precedence over auto SDPs.

If a manual SDP is available, it should be used to create the auto SDP binding; otherwise, the PE router creates auto SDPs and auto SDP bindings. If there are multiple SDPs available for a remote PE router, the SDP with the highest SDP ID should be used.

If a manual SDP was provisioned after the RADIUS discovery process, the **admin>radius-discovery>force-discover** command should be executed in order to use the new manual SDPs.

Users cannot remove a manual SDP if there us an auto SDP binding associated with the manual SDP. However, manual SDP bindings and RADIUS discovery are mutually exclusive. If a manual SDP binding was already provisioned for a service, RADIUS discovery cannot be enabled for the service, If RADIUS discovery was already enabled for a service, manual SDP bindings cannot be provisioned for the service.

Discovery Procedures

A PE router issues an access-request to the RADIUS server using the configured VPN as the user-name and configured password. The service type of the access request is L2VPN.4

The RADIUS server authenticates the PE router. If authentication is successful, it responds with an access-accept which includes a list of IP addresses of the PE routers. Optional information such as vc-type and SDP could be included. If authentication fails, an access-reject message is returned.

The access-accept response has a session-timeout attribute in which a PE router needs to issue a new access-request before the access-accept times out.

Auto SDPs and SDP-bindings Creation

After the discovery of all other PE routers, the PE router checks if it has valid auto SDP bindings for remote PE routers. Manual SDPs always take precedence over auto SDPs. If there was a manual SDP available, it is used to create the auto SDP binding; otherwise, the PE router creates auto SDPs and auto SDP bindings. If there were multiple SDPs available for a PE router, the SDP with the smallest SDP ID is used.

What auto-SDP to create, LDP or GRE, depends on the RADIUS configuration, default is GRE.

When an auto SDP is created for a remote PE router and there is no targeted LDP session to the PE router, a targeted LDP session is automatically created.

Pseudo-Wire Setup

The PE routers use “AII” as the VC-ID (PW ID) to signal pseudowires to each other by targeted LDP.

RADIUS Server Polling

A PE router should periodically query the RADIUS server to make sure its L2VPN membership information is up-to-date. The polling interval is configurable by CLI.

Any change to the VPN membership (such as adding or removing a PE) takes effect at the next polling.

Removing RADIUS Discovery

The VPLS service must be shut down in order to disable or remove RADIUS discovery from a VPLS service.

When RADIUS discovery is disabled for a VPLS service, auto SDPs and auto SDP-bindings for that VPLS are removed.

Remove a PE from VPLS

When a PE is removed from a VPLS service, the RADIUS server must be updated manually to remove the PE from the membership database.

VPLS Service Considerations

This section describes various of the general 7750 SR service features and any special capabilities or considerations as they relate to VPLS services.

SAP Encapsulations

The 7750 SR VPLS service is designed to carry Ethernet frame payloads, so it can provide connectivity between any SAPs and SDPs that pass Ethernet frames. The following SAP encapsulations are supported on the 7750 SR VPLS service:

- Ethernet null
- Ethernet dot1q
- Ethernet qinq
- SONET/SDH BCP-null
- SONET/SDH BCP-dot1q
- ATM VC with RFC 2684 Ethernet bridged encapsulation (See ATM/Frame Relay PVC Access and Termination on a VPLS Service on page 426.)
- FR VC with RFC 2427 Ethernet bridged encapsulation (See ATM/Frame Relay PVC Access and Termination on a VPLS Service on page 426.)

VLAN processing

The SAP encapsulation definition on Ethernet ingress ports defines which VLAN tags are used to determine the service that the packet belongs to:

1. **null** encap defined on ingress: Any VLAN tags are ignored and the packet goes to a default service for the SAP;
2. **dot1q** encap defined on ingress: Only first label is considered;
3. **qinq** encap defined on ingress: both labels are considered.
   Note that the SAP can be defined with a wildcard for the inner label. (e.g. “100:*”). In this situation all packets with an outer label of 100 will treated as belonging to the SAP. If, on the same physical link there is also a SAP defined with qinq encap of 100:1 then traffic with 100:1 will go to that SAP and all other traffic with 100 as the first label will go to the SAP with the 100:* definition.

In situations 2 and 3 above, traffic encapsulated with tags for which there is no definition are discarded.

Configuring a VPLS Service with CLI

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

Topics in this section include:

List of Commands

Table 13 lists all the service configuration commands indicating the configuration level at which each command is implemented with a short command description. VPLS services are configured in the `config>service>vpls` context. The command list is organized in the following task-oriented manner:

- Configure a VPLS service
  - → Configure VPLS STP parameters
  - → Configure VPLS multicast FIB parameters
  - → Configure VPLS IGMP parameters
  - → Configure Split Horizon Group parameters
- Configure VPLS multicast FIB parameters
  - → Configure VPLS IGMP parameters
- Configure a VPLS SAP
  - → Configure SAP ATM parameters
  - → Configure SAP egress parameters
  - → Configure SAP ingress parameters
  - → Configure SAP STP parameters
  - → Configure SAP subscriber management parameters
- Configure mesh SDP parameters
- Configure spoke SDP parameters
  - → Configure spoke SDP STP parameters
- Configure an egress multicast group
  - → Applying an egress multicast group to a VPLS SAP
- Configure a PE discovery policy
  - → Applying a PE discovery policy to VPLS

Table 1: CLI Commands to Configure VPLS Service Parameters

Command	Description	Page
Configure a VPLS service		
`config>service# vpls` *`service-id`* `[customer` *`customer-id`*`] [vpn` *`vpn-id`*`]`		459
`customer`	Specifies the existing customer ID number to be associated with the service.	459
`vpn`	Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.	459

Table 1: CLI Commands to Configure VPLS Service Parameters (Continued)

Command	Description	Page
def-mesh-vc-id	Configures the value used by each end of a tunnel to identify the VC.	460
description	A text string describing the service.	458
disable-aging	Disables MAC address aging across a VPLS service.	461
disable-learning	Disables MAC address learning across a VPLS service.	461
discard-unknown	Discards unknown destination addresses.	462
fdb-table-high-wmark	Specifies the value to send logs and traps when the threshold is reached.	462
fdb-table-low-wmark	Configures the low watermark for the FDB table.	462
fdb-table-size	Specifies the maximum number of MAC entries in the forwarding database (FDB).	462
local-age	Specifies the aging time for locally learned MAC addresses in the FDB.	466
mac-move	Enables the context to configure MAC move attributes.	466
mac-protect	Indicates whether or not this MAC is protected.	467
move-frequency	Indicates the maximum rate at which MAC's can be re-learned in the VPLS service, before the SAP where the moving MAC was last seen is automatically disabled in order to protect the system against undetected loops or duplicate MAC's.	468
retry-timeout	Indicates the time in seconds to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled.	468
remote-age	Specifies the aging time for remotely learned MAC addresses in the FDB.	469
send-flush-on-failure	Enables sending MAC withdraw message on SAP/spoke-SDP failure.	470
service-mtu	Configures the service payload (Maximum Transmission Unit – MTU) in bytes for the service ID overriding the service-type default MTU.	470
no shutdown	Administratively enables the service.	457
Configure VPLS STP parameters		
config>service# vpls		
stp	Context for configuring the STP bridge-level parameters for the VPLS instance.	483
forward-delay	Configures STP forward delay timer for the VPLS instance.	484
hello-time	Configures STP hello time for the VPLS instance.	484
hold-count	Configure BPDU transmit hold count.	485
max-age	Configures STP max age for the VPLS instance.	487
mst-instance	Creates the context to configure MST instance (MSTI) related parameters.	487
mst-priority	Specifies the bridge priority for this specific Multiple Spanning Tree instance for this service.	488

Table 1: CLI Commands to Configure VPLS Service Parameters (Continued)

Command	Description	Page
vlan-range	Specifies a range of VLANs associated with a certain MST-instance. This range applies to all SAPs of the mVPLS.	488
mst-max-hops	Specifies the number of hops in the region before BPDU is discarded and the information held for the port is aged out.	489
mst-revision	Specifies the MSTP region revision number. Together with region name and VLAN-to-instance assignment it defines the MSTP region.	489
mst-name	Defines an MST region name.	489
mode	Configure protocol version.	487
priority	Configures STP bridge priority for the VPLS instance. When running MSTP, this is the bridge priority used for the CIST.	491
no shutdown	Administratively enables the SAP STP.	457
Configure VPLS multicast FIB parameters		
config>service>vpls#		
mfib-table-high-wmark	Configures the high watermark for the multicast FIB table.	468
mfib-table-low-wmark	Configures the low watermark for the multicast FIB table.	469
mfib-table-size	Configures the maximum number of SG entries in the multicast FIB table.	469
Configure VPLS IGMP parameters		
config>service>vpls# igmp-snooping		
query-interval	Configures the IGMP query interval.	539
report-src-ip	Specifies the source IP address used when generating IGMP reports.	540
robust-count	Configures the IGMP robustness variable.	541
config>service>vpls>sap# igmp-snooping		
config>service>vpls>mesh-sdp# igmp-snooping		
config>service>vpls>spoke-sdp# igmp-snooping		
fast-leave	Enables IGMP fast-leave processing.	534
import	Imports a policy to filter IGMP packets.	535
last-member-query-interval	Configures the IGMP last member query interval.	536
max-num-groups	Configures the max number of multicast groups allowed.	538
mrouter-port	For SAP and spoke SDP configurations, specifies whether a multicast router is attached behind this SAP or SDP.	539
mvr	Enables the context to configure Multicast VPLS Registration (MVR) parameters.	539

Table 1: CLI Commands to Configure VPLS Service Parameters (Continued)

Command	Description	Page
group-policy	Configures the policy containing channels to import on the Multicast VPLS Registration (MVR) VPLS.	535
from-vpls	Configures the VPLS from which to import channels.	534
to-sap	Configures the SAP to which to copy the channels imported with MVR.	543
query-interval	Configures the IGMP query interval.	539
query-response-interval	Configures the IGMP query response interval.	540
robust-count	Configures the IGMP robustness variable.	541
send-queries	Enables generation of IGMP general queries.	541
static	Enables access to the context to configure static group addresses.	542
group	Adds a static multicast group as a (*, g) or as one or more (s,g) records.	534
source	Adds a statstargic (s,g) entry to allow multicast traffic for the corresponding multicast group from that specific source.	541
starg	Adds a static (*,g) entry to allow multicast traffic for the corresponding multicast group from any source.	542
version	Specifies the version of IGMP which is running on this SDP.	542
Configure Split Horizon Group parameters		
config>service>vpls>split-horizon-group		
config>service>vpls>sap		
restrict-protected-src	Indicates how the agent will handle relearn requests for protected MAC addresses.	472
restrict-unprotected-dst	Indicates how the system will forward packets destined to an unprotected MAC address. When enabled, packets destined to an unprotected MAC address will be dropped.	472
Configure a VPLS SAP		
config>service>vpls# sap *sap-id* [split-horizon-group *group-name*]		
accounting-policy	Configures the accounting policy that applies to the SAP.	522
anti-spoof	Enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP.	545
arp-reply-agent	Enables a special ARP response mechanism in the system for ARP requests destined to static or dynamic hosts associated with the SAP.	545
authentication-policy	Defines which subscriber authentication policy must be applied when a DHCP message is received on the interface.	522
bpdu-translation	Enables the translation of BPDUs to a given format, meaning that all BPDUs transmitted on a given SAP will have a specified format.	460

Table 1: CLI Commands to Configure VPLS Service Parameters (Continued)

Command	Description	Page
collect-stats	Enables the collection of accounting and statistical data for the SAP or network port.	522
description	A text string describing the SAP.	
disable-aging	Disables aging of MAC addresses for this SAP.	461
disable-learning	Disables learning of new MAC addresses for this SAP.	461
discard-unknown-source	Enables the discard of packets with unknown source MAC addresses when the FDB limit is reached.	496
host	Creates a static subscriber host for the SAP. Applications within the system that make use of static host entries include anti-spoof, ARP reply agent and source MAC population into the VPLS forwarding database.	546
host-connectivity-verify	Enables subscriber host connectivity verification on a given SAP within a VPLS service.	547
l2pt-termination	Enables L2PT termination on a given SAP or spoke SDP. L2TP termination will be supported only for STP BPDUs. PDUs of other protocols will be discarded.	460
limit-mac-move	Indicates whether or not the agent will limit the MAC re-learn (move) rate on this SAP.	497
mac-pinning	Disables relearning MAC addresses on other SAPs.	497
managed-vlan-list	Enables the context to provision a VLAN list for the managed VPLS SAP.	499
max-nbr-mac-addr	Configure the maximum number of MAC entries in the FDB.	498
multi-service-site	Specifies the multi-service-site to which this SAP belongs.	498
restrict-protected-src	Indicates how the agent will handle relearn requests for protected MAC addresses.	472
restrict-unprotected-dst	Indicates how the system will forward packets destined to an unprotected MAC address.	472
static-mac	Creates a local static MAC entry in the Virtual Private LAN Service (VPLS) forwarding database (FDB) associated with the Service Access Point (SAP).	499
Configure GSMP parameters		
config>service>vpls>gsmp		
group	Specifies a GSMP name. A GSMP group name is unique only within the scope of the service in which it is defined.	474
ancp	Configures ANCP parameters for this GSMP group.	474
dynamic-topology-discover	Enables the ANCP dynamic topology discovery capability.	474
line-configuration	Enables the ANCP line-configuration capability.	474

Table 1: CLI Commands to Configure VPLS Service Parameters (Continued)

Command	Description	Page
oam	Specifies whether or not the GSMP ANCP OAM capability should be negotiated at startup of the GSMP connection.	475
hold-multiplier	Configures the hold-multiplier for the GSMP connections in this group.	475
keepalive	Configures keepalive values for the GSMP connections in this group.	475
neighbor	Configures a GSMP ANCP neighbor.	475
local-address	Configures the source ip-address used in the connection towards the neighbor.	476
priority-marking	Configures the type of priority marking to be used.	476
no shutdown	Administratively disables a GSMP entity.	457
Configure SAP ATM parameters		
config>service>vpls>sap>atm#		
egress	Enables the context to configure egress ATM attributes for the SAP.	501
ingress	Enables the context to configure ingress ATM attributes for the SAP.	502
traffic-desc	Assigns an ATM traffic descriptor profile to a given context.	502
encapsulation	Configures RFC2684 encapsulation for an ATM PVCC delimited SAP.	501
oam	Enables the context to configure OAM functionality for a PVCC delimiting a SAP.	502
alarm-cells	Configures AIS/RDI fault management on a PVCC.	503
Configure SAP egress parameters		
config>service>vpls>sap#		
egress	Configures egress SAP QoS policies and filter policies.	504
filter	Associates an IP filter policy or MAC filter policy with an egress SAP.	505
qos	Associates a Quality of Service (QoS) policy with an egress SAP.	507
queue-override	Enables the context to configure override values for the specified SAP egress QoS queue.	508
queue	Specifies the ID of the queue whose parameters are to be overridden.	508
adaptation-rule	Overrides specific attributes of the specified queue's adaptation rule parameters.	509
cbs	Overrides specific attributes of the specified queue's CBS parameters.	511
high-prio-only	Overrides specific attributes of the specified queue's high-prio-only parameters.	512
mbs	Overrides specific attributes of the specified queue's MBS parameters.	512
rate	Overrides specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.	514

Table 1: CLI Commands to Configure VPLS Service Parameters (Continued)

Command	Description	Page
scheduler-override	Specifies the set of attributes whose values have been overridden via management on this virtual scheduler.	515
scheduler-policy	Associates an existing scheduler policy to an egress scheduler used by SAP queues associated with this multi-service customer site.	517
Configure SAP ingress parameters		
config>service>vpls>sap#		
ingress	Configures ingress SAP QoS policies and filter policies.	504
filter	Associates an IP filter policy or MAC filter policy with an ingress SAP or IP interface.	505
match-qinq-dot1p	Configures filtering based on the p-bits in the top or bottom tag of a Q-in-Q encapsulated Ethernet frame.	518
qos	Associates a Quality of Service (QoS) policy with an ingress SAP or IP interface.	507
queue-override	Enables the context to configure override values for the specified SAP egress QoS queue.	508
queue	Specifies the ID of the queue whose parameters are to be overridden.	508
adaptation-rule	Overrides specific attributes of the specified queue's adaptation rule parameters.	509
cbs	Overrides specific attributes of the specified queue's CBS parameters.	511
high-prio-only	Overrides specific attributes of the specified queue's high-prio-only parameters.	512
mbs	Overrides specific attributes of the specified queue's MBS parameters.	513
rate	Overrides specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.	514
scheduler-override	Specifies the set of attributes whose values have been overridden via management on this virtual scheduler.	515
scheduler-policy	Associates an existing scheduler policy to an egress scheduler used by SAP queues associated with this multi-service customer site.	517
scheduler-policy	Associates an existing scheduler policy to an ingress scheduler used by SAP queues associated with this multi-service customer site.	517
Configure SAP STP parameters		
config>service>vpls>sap>stp#		
auto-edge	Configures automatic detection of the edge port characteristics of the SAP or spoke SDP.	483
edge-port	Configures the SAP as an edge or non-edge port.	483
link-type	Instructs STP on the maximum number of bridges behind this SAP.	485

Table 1: CLI Commands to Configure VPLS Service Parameters (Continued)

Command	Description	Page
mst-instance	Enables the context to configure MSTI related parameters at SAP level.	486
mst-path-cost	Specifies path-cost within a given instance, expressing probability that given port will be put into forwarding state in case loop occurs.	486
mst-port-priority	Specifies the port priority within a given instance, expressing probability that given port will be put into a forwarding state if loop occurs.	486
path-cost	Configures the Spanning Tree Protocol (STP) path cost for the SAP. These are the values used for CIST when running MSTP.	490
port-num	Configures the virtual port number which uniquely identifies a SAP within configuration bridge protocol data units (BPDUs).	490
priority	Configures the Spanning Tree Protocol (STP) priority for the SAP. These are the values used for CIST when running MSTP.	491
root-guard	Specifies whether this port is allowed to become an STP root port.	529
no shutdown	Administratively enables an entity.	457
Configure SAP subscriber management parameters		
config>service>vpls>sap>sub-sla-mgmt		
def-sla-profile	Specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sla-profile** context.	530
def-sub-profile	Specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-profile** context.	530
mac-da-hashing	Specifies whether subscriber traffic egressing a LAG SAP has its egress LAG link selected by a function of the MAC destination address instead of the subscriber ID.	531
multi-sub-sap	Configures the maximum number of subscribers for this SAP.	531
single-sub-parameters	Enables enables the context to configure single subscriber parameters for this SAP.	532
non-sub-traffic	Configures non-subscriber traffic profiles.	531
profiled-traffic-only	Enables profiled traffic only for this SAP.	532
sub-ident-policy	Associates a subscriber identification policy to this SAP. The subscriber identification policy must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-ident-policy** context.	533
no shutdown	Administratively enables an entity.	457
Configure mesh SDP parameters		
config>service>vpls# mesh-sdp *sdp-id*[:*vc-id*] [vc-type {ether|vlan}]		
accounting-policy	Configures the accounting policy that applies to the SDP.	522

Table 1: CLI Commands to Configure VPLS Service Parameters (Continued)

Command	Description	Page
collect-stats	Enables the collection of accounting and statistical data for the SDP.	522
egress	Configures the egress label.	526
ingress	Configures the ingress label.	526
filter	Associates an IP filter policy or MAC filter policy with an egress Service Access Point (SAP) or IP interface.	505
mfib-allowed-mda-destinations	Enables the context to configure MFIB-allowed MDA destinations.	526
mda	Specifies an MFIB-allowed MDA destination for an SDP binding configured in the system.	527
mac-pinning	Disables relearning MAC addresses on other SAPs	497
static-mac	Creates a remote static MAC entry in the VPLS FDB associated with the SDP.	528
vc-label	Configures the egress or ingress VC label.	527
vlan-vc-tag	Configures the VLAN VC tag.	529
Configure spoke SDP parameters		
config>service>vpls# spoke-sdp *sdp-id* [vc-type {ether|vlan}] [split-horizon-group *group-name*]		
accounting-policy	Configures the accounting policy that applies to the SDP.	522
bpdu-translation	Enables the translation of BPDUs to a given format, meaning that all BPDUs transmitted on a given SDP will have a specified format.	460
collect-stats	Enables the collection of accounting and statistical data for the SDP.	522
disable-aging	Disables aging of MAC addresses for this SDP.	461
disable-learning	Disables learning of new MAC addresses for this SDP.	461
discard-unknown-source	Enables the discard of packets with unknown destination MAC addresses.	496
egress	Configures the egress label.	526
ingress	Configures the ingress label.	526
filter	Associates an IP filter policy or MAC filter policy with an egress Service Access Point (SAP) or IP interface.	505
mfib-allowed-mda-destinations	Enables the context to configure MFIB-allowed MDA destinations.	526
vc-label	Configures the egress VC label.	527
l2pt-termination	Enables L2PT termination on a given SAP or spoke SDP. L2TP termination will be supported only for STP BPDUs. PDUs of other protocols will be discarded.	460
limit-mac-move	Indicates whether or not the agent will limit the MAC re-learn (move) rate on this SDP.	497

Table 1: CLI Commands to Configure VPLS Service Parameters (Continued)

Command	Description	Page
max-nbr-mac-addr	Configure the max number of MAC entries in the FDB for this SDP.	498
static-mac	Creates a remote static MAC entry in the VPLS FDB associated with the SDP.	499
vlan-vc-tag	Configures the VLAN VC tag.	529
no shutdown	Administratively enables the SDP.	457
Configure spoke SDP STP parameters		
config>service>vpls>spoke-sdp>stp		
auto-edge	Enables automatic detection of edge port.	483
edge-port	Configure SAP as edge or non-edge port.	483
link-type	Configure link type of the SAP.	485
path-cost	Configures the STP path cost for the spoke SDP.	490
port-num	Configures virtual port number.	490
priority	Configures the STP priority for the spoke SDP.	491
root-guard	Specifies whether this port is allowed to become an STP root port.	529
no shutdown	Administratively enables STP for the spoke SDP.	457
Configure an egress multicast group		
config>service>egress-multicast-group		
description	Defines an ASCII string associated with the egress multicast group.	551
dest-chain-limit	Defines the maximum length of an egress forwarding plane efficient multicast replication chain for an egress-multicast-group.	551
sap-common-requirements	Configures the common SAP parameter requirements. The SAP common requirements are used to evaluate each SAP for group membership.	552
dot1q-etype	Specifies the dot1q EtherType that must exist on the SAP's access port to allow the SAP membership within the egress-multicast-group.	554
egress-filter	Identifies the type of filter and actual filter ID that must be provisioned on the SAP prior to the SAP being made a member of the egress-multicast-group.	553
encap-type	Specifies the encapsulation type that must exist on the SAP's access port to allow the SAP membership within the egress-multicast-group.	553
Applying an egress multicast group to a VPLS SAP		
config>service>vpls>sap>egress		
multicast-group	Assigns a VPLS Ethernet SAP into an egress multicast group.	506
Configure a PE discovery policy		

Table 1: CLI Commands to Configure VPLS Service Parameters (Continued)

Command	Description	Page
`config>service>pe-discovery-policy` *name*		
`password`	Configures the PE discovery password that is used when contacting the RADIUS server for VPLS auto-discovery.	556
`polling-interval`	Configures the PE discovery polling interval.	556
`server`	Identifies a specific RADIUS server.	557
`timeout`	Configures the time to wait before timing out a RADIUS server.	557
Applying a PE discovery policy to VPLS		
`config>service>vpls>radius-discovery`		
`radius-discovery`	Enables the RADIUS provider edge discovery for this VPLS service.	558
`pe-discovery-policy`	Specifies the existing RADIUS PE discovery policy name.	558
`user-name-format`	Specifies whether the RADIUS user name is a VPN ID or router-distinguisher.	558

Basic Configuration

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

- Customer ID (refer to Configuring Customers on page 64)
- For a local service, configure two SAPs, specifying local access ports and encapsulation values.
- For a distributed service, configure a SAP and an SDP for each far-end node.

The following example displays a sample configuration of a local VPLS service on ALA-1.

```
*A:ALA-1>config>service>vpls# info
----------------------------------------------
...
        vpls 9001 customer 6 create
            description "Local VPLS"
            stp
                shutdown
            exit
            sap 1/2/2:0 create
                description "SAP for local service"
            exit
            sap 1/1/5:0 create
                description "SAP for local service"
            exit
            no shutdown
----------------------------------------------
*A:ALA-1>config>service>vpls#
```

The following example displays a sample configuration of a distributed VPLS service between ALA-1, ALA-2, and ALA-3.

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vpls 9000 customer 6 create
            shutdown
            description "This is a distributed VPLS."
            def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/1/5:16 create
                description "VPLS SAP"
            exit
            spoke-sdp 2:22 create
            exit
            mesh-sdp 7:750 create
            exit
        exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

```
*A:ALA-2>config>service# info
----------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
            def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/1/5:16 create
                description "VPLS SAP"
            exit
            spoke-sdp 2:22 create
            exit
            mesh-sdp 8:750 create
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-2>config>service#


*A:ALA-3>config>service# info
----------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
           def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/1/3:33 create
                description "VPLS SAP"
            exit
            spoke-sdp 2:22 create
            exit
            mesh-sdp 8:750 create
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-3>config>service#
```

Common Configuration Tasks

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

For egress multicast groups (optional):

1. Define egress multicast group name(s)
2. Specify the destinations per pass
3. Define SAP common requirements

For VPLS services:

1. Associate VPLS service with a customer ID
2. Define SAPs:

   — Select node(s) and port(s)

   — Optional - select QoS policies other than the default (configured in `config>qos` context)

   — Optional - select filter policies (configured in `config>filter` context)

   — Optional - select accounting policy (configured in `config>log` context)

   — Optional - specify SAP egress multicast-group name
3. Associate SDPs for (distributed services)
4. Modify STP default parameters (optional) (see VPLS and Spanning Tree Protocol on page 335)
5. Enable service

Configuring VPLS Components

Use the CLI syntax displayed below to configure the following entities:

Configuring Egress Multicast Groups

Use the following CLI syntax to create a VPLS service:

CLI Syntax:

```
config>service# egress-multicast-group group-name
                description description-string
                dest-chain-limit [destinations per pass]
                sap-common-requirements
                   dot1q-etype 0x0600..0xffff
                   egress-filter [ip ip-filter-id]
                   egress-filter [ipv6 ipv6-filter-id]
                   egress-filter [mac mac-filter-id]
                   no egress-filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                      [mac mac-filter-id]
                   encap-type {dot1q|null}
```

Example:

```
config>service# egress-multicast-group "vpls-emg-1"
      config>service>egress-multicast-group# dest-chain-limit 10
      config>service>egress-multicast-group# sap-common-requirements
      config>service>egress-multicast-group>sap-common-requirements#
dot1q-etype 0x060e
      config>service>egress-multicast-group>sap-common-requirements#
```

```
egress-filter ip 10
       config>service>egress-multicast-group>sap-common-requirements#
exit
```

The following example displays the egress multicast group configuration:

```
A:ALA-49>config>service>egress-multicast-group# info
----------------------------------------------
                dest-chain-limit 10
                sap-common-requirements
                    dot1q-etype 0x060e
                    egress-filter ip 10
                exit
----------------------------------------------
A:ALA-49>config>service>egress-multicast-group#
```

Creating a VPLS Service

Use the following CLI syntax to create a VPLS service:

CLI Syntax: config>service# vpls *service-id* [customer *customer-id*] [vpn *vpn-id*] [m-vpls]
description *description-string*
no shutdown

Example:
```
config>service# vpls 9000 customer 6 create
config>service>vpls$ description "This is a distributed VPLS."
config>service>vpls# def-mesh-vc-id 750
config>service>vpls# no shutdown
config>service>vpls# exit
```

The following example displays the VPLS configuration:

```
*A:ALA-1>config>service>vpls# info
----------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
            def-mesh-vc-id 750
            stp
                shutdown
            exit
        exit
...
----------------------------------------------
*A:ALA-1>config>service>vpls#
```

Enabling MAC Move

The mac-move feature is useful to protect against undetected loops in your VPLS topology as well as the presence of duplicate MACs in a VPLS service. For example, if two clients in the VPLS have the same MAC address, the VPLS will experience a high re-learn rate for the MAC and will shut down the SAP or spoke SDP when the threshold is exceeded.

Use the following CLI syntax to configure **mac-move** parameters.

CLI Syntax: config>service# vpls *service-id* [customer *customer-id*] [vpn *vpn-id*] [m-vpls]
mac-move
move-frequency *frequency*
retry-timeout *timeout*
no shutdown

Example:
```
config>service# vpls 9000 customer 6
config>service>vpls# mac-move
config>service>vpls>mac-move# retry-timeout 30
config>service>vpls>mac-move# move-frequency 5
```

```
config>service>vpls>mac-move# no shutdown
config>service>vpls>mac-move# exit
```

The following example displays the mac-move configuration:

```
A:ALA-48>config>service# info
----------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
            def-mesh-vc-id 750
            stp
                no shutdown
            exit
            mac-move
                move-frequency 5
                retry-timeout 30
                no shutdown
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-48>config>service#
```

Configuring STP Bridge Parameters in a VPLS

Modifying some of the Spanning Tree Protocol parameters allows the operator to balance STP between resiliency and speed of convergence extremes. Modifying particular parameters, mentioned below, must be done in the constraints of the following two formulae:

2 x (*Bridge_Forward_Delay* - 1.0 seconds) >= *Bridge_Max_Age*

Bridge_Max_Age >= 2 x (*Bridge_Hello0_Time* + 1.0 seconds)

The following STP parameters can be modified at VPLS level:

- Bridge STP Admin State on page 380
- Mode on page 381
- Bridge Priority on page 381
- Max Age on page 381
- Forward Delay on page 382
- Hello Time on page 383
- MST Instances on page 384
- MST Max Hops on page 384
- MST Name on page 384
- MST Revision on page 384

STP always uses the locally configured values for the first three parameters (Admin State, Mode and Priority).

For the parameters Max Age, Forward Delay, Hello Time and Hold Count, the locally configured values are only used when this bridge has been elected root bridge in the STP domain, otherwise the values received from the root bridge are used. The exception to this rule is: when STP is running in RSTP mode, the Hello Time is always taken from the locally configured parameter. The other parameters are only used when running mode MSTP.

Bridge STP Admin State

The administrative state of STP at the VPLS level is controlled by the shutdown command.

When STP on the VPLS is administratively disabled, any BPDUs are forwarded transparently through the 7750 SR. When STP on the VPLS is administratively enabled, but the administrative state of a SAP or spoke SDP is down, BPDUs received on such a SAP or spoke SDP are discarded.

CLI Syntax: `config>service>vpls` *`service-id`*`# stp`
`no shutdown`

Mode

To be compatible with the different iterations of the IEEE 802.1D standard, the 7750 SR supports several variants of the Spanning Tree protocol:

- `rstp` — Rapid Spanning Tree Protocol (RSTP) compliant with IEEE 802.1D-2004 - default mode
- `dot1w` — Compliant with IEEE 802.1w
- `comp-dot1w` — Operation as in RSTP but backwards compatible with IEEE 802.1w (this mode was introduced for interoperability with some MTU types)
- `mstp` — Compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q REV/D5.0-09/2005 . This mode of operation is only supported in an mVPLS.

See section Spanning Tree Operating Modes on page 335 for details on these modes.

CLI Syntax: `config>service>vpls` *`service-id`*`# stp`
`mode {rstp | comp-dot1w | dot1w | mstp}`
Default: `rstp`

Bridge Priority

The `bridge-priority` command is used to populate the priority portion of the bridge ID field within outbound BPDUs (the most significant 4 bits of the bridge ID). It is also used as part of the decision process when determining the best BPDU between messages received and sent. When running MSTP, this is the bridge priority used for the CIST.

All values will be truncated to multiples of 4096, conforming with IEEE 802.1t and 802.1D-2004.

CLI Syntax: `config>service>vpls` *`service-id`*`# stp`
`priority` *`bridge-priority`*
Range: 1 to 65535
Default: 32768
Restore Default: `no priority`

Max Age

The `max-age` command indicates how many hops a BPDU can traverse the network starting from the root bridge. The message age field in a BPDU transmitted by the root bridge is initialized to 0. Each other bridge will take the message_age value from BPDUs received on their root port and increment this value by 1. The message_age thus reflects the distance from the root bridge. BPDUs with a message age exceeding max-age are ignored.

STP uses the max-age value configured in the root bridge. This value is propagated to the other bridges via the BPDUs.

The default value of `max-age` is 20. This parameter can be modified within a range of 6 to 40, limited by the standard STP parameter interaction formulae.

CLI Syntax: `config>service>vpls` *`service-id`*`# stp`
`max-age` *`max-info-age`*
Range: 6 to 40 seconds
Default: 20 seconds
Restore Default: `no max-age`

Forward Delay

RSTP, as defined in the IEEE 802.1D-2004 standards, will normally transition to the forwarding state via a handshaking mechanism (rapid transition), without any waiting times. If handshaking fails (e.g. on shared links, see below), the system falls back to the timer-based mechanism defined in the original STP (802.1D-1998) standard.

A shared link is a link with more than two Ethernet bridges (for example, a shared 10/100BaseT segment). The `port-type` command is used to configure a link as point-to-point or shared (see section SAP Link Type on page 392).

For timer-based transitions, the 802.1D-2004 standard defines an internal variable forward-delay, which is used in calculating the default number of seconds that a SAP or spoke SDP spends in the discarding and learning states when transitioning to the forwarding state.

The value of the forward-delay variable depends on the STP operating mode of the VPLS instance:

- in `rstp` mode, but only when the SAP or spoke SDP has not fallen back to legacy STP operation, the value configured by the `hello-time` command is used;
- in all other situations, the value configured by the `forward-delay` command is used.

CLI Syntax: `config>service>vpls` *`service-id`*`# stp`
`forward-delay` *`seconds`*
Range: 4 to 30 seconds
Default: 15 seconds
Restore Default: `no forward-delay`

Hello Time

The `hello-time` command configures the Spanning Tree Protocol (STP) hello time for the Virtual Private LAN Service (VPLS) STP instance.

The *seconds* parameter defines the default timer value that controls the sending interval between BPDU configuration messages by this bridge, on ports where this bridge assumes the designated role.

The active hello time for the spanning tree is determined by the root bridge (except when the STP is running in RSTP mode, then the hello time is always taken from the locally configured parameter).

The configured hello-time value can also be used to calculate the bridge forward delay, see Forward Delay on page 382

CLI Syntax: `config>service>vpls` *`service-id`*`# stp`
`hello-time` *`hello-time`*
Range: 1 to 10 seconds
Default: 2 seconds
Restore Default: `no hello-time`

Hold Count

The `hold-count` command configures the peak number of BPDUs that can be transmitted in a period of one second.

CLI Syntax: `config>service>vpls` *`service-id`*`# stp`
`hold-count` *`count-value`*
Range: 1 to 10
Default: 6
Restore Default: `no hold-count`

MST Instances

You can create up to 15 MST-instances. They can range from 1 to 4094. By changing path-cost and priorities, you can make sure that each instance will form it's own tree withing the region, thus making sure different VLANs follow different paths.

You can assign non overlapping VLAN ranges to each instance. VLANs that are not assigned to an instance are implicitly assumed to be in instance 0, which is also called the CIST. This CIST cannot be deleted or created.

The parameter that can be defined per instance are mst-priority and vlan-range.

- mst-priority — The bridge-priority for this specific mst-instance. It follows the same rules as bridge-priority. For the CIST, the bridge-priority is used.
- vlan-range — The VLANs are mapped to this specific mst-instance. If no VLAN-ranges are defined in any mst-instances, then all VLANs are mapped to the CIST.

MST Max Hops

The mst-max-hops command defines the maximum number of hops the BPDU can traverse inside the region. Outside the region max-age is used.

MST Name

The MST name defines the name that the operator gives to a region. Together with MST revision and the VLAN to MST-instance mapping, it forms the MST configuration identifier. Two bridges that have the same MST configuration identifier form a region if they exchange BPDUs.

MST Revision

The MST revision together with MST-name and VLAN to MST-instance mapping define the MST configuration identifier. Two bridges that have the same MST configuration identifier form a region if they exchange BPDUs.

Configuring GSMP Parameters

The following parameters must be configured in order for GSMP to function:

- One or more GSMP sessions
- One or more ANCP policies
- For basic subscriber management only, ANCP static maps
- For enhanced subscriber management only, associate subscriber profiles with ANCP policies.

This example configures a GSMP group.

Example: config>service>vpls# gsmp
```
    config>service>vpls>gsmp>group# description "test group config"
    config>service>vpls>gsmp# group "group1" create
    config>service>vpls>gsmp>group# ancp
    config>service>vpls>gsmp>group# description "test group config"
    config>service>vpls>gsmp>group# neighbor 10.10.10.104 create
    config>service>vpls>gsmp>group>neighbor# description "neighbor1
config
    config>service>vpls>gsmp>group>neighbor# local-address 10.10.10.103
    config>service>vpls>gsmp>group>neighbor# no shutdown
    config>service>vpls>gsmp>group>neighbor# exit
    config>service>vpls>gsmp>group# exit
    config>service>vpls>gsmp# no shutdown
```

This example displays the GSMP group configuration.

```
A:ALA-48>config>service>vpls>gsmp# info
----------------------------------------------
                group "group1" create
                    description "test group config"
                    neighbor 10.10.10.104 create
                        description "neighbor1 config"
                        local-address 10.10.10.103
                        no shutdown
                    exit
                    no shutdown
                exit
                no shutdown
----------------------------------------------
A:ALA-48>config>service>vpls>gsmp#
```

Configuring a VPLS SAP

A default QoS policy is applied to each ingress and egress SAP. Additional QoS policies can be configured in the `config>qos` context. There are no default filter policies. Filter policies are configured in the `config>filter` context and must be explicitly applied to a SAP.

Use the following CLI syntax to create:

- Local VPLS SAPs on page 386
- Distributed VPLS SAPs on page 387

Local VPLS SAPs

To configure a local VPLS service, enter the `sap` *`sap-id`* command twice with different port IDs in the same service configuration.

Example:
```
config>service# vpls 9001 customer 6 create
          config>service>vpls# sap 1/2/2:0 create
          config>service>vpls>sap$ no shutdown
          config>service>vpls>sap# description "SAP for local service"
          config>service>vpls>sap# exit
          config>service>vpls# sap 1/1/5:0 create
          config>service>vpls>sap$ no shutdown
          config>service>vpls>sap# description "SAP for local service"
          config>service>vpls>sap# exit
```

The following example displays the local VPLS configuration:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vpls 90001 customer 6 create
            description "Local VPLS"
            stp
                shutdown
            exit
            sap 1/2/2:0 create
                description "SAP for local service"
            exit
            sap 1/1/5:0 create
                description "SAP for local service"
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-1>config>service#
```

Distributed VPLS SAPs

To configure a distributed VPLS service, you must configure service entities on originating and far-end nodes. You must use the same service ID on all ends (for example, create a VPLS service ID 9000 on ALA-1, ALA-2, and ALA-3). A distributed VPLS consists of a SAP on each participating node and an SDP bound to each participating node.

For SDP configuration information, see Configuring an SDP on page 68. For SDP binding information, see Configuring SDP Bindings on page 400.

This example configures a distributed service between ALA-1, ALA-2, and ALA-3.

Example:

```
*A:ALA-1>config>service# vpls 9000 customer 6 create
*A:ALA-1>config>service>vpls# sap 1/2/5:0 create
*A:ALA-1>config>service>vpls>sap$ no shutdown
*A:ALA-1>config>service>vpls>sap# multi-service-site west
*A:ALA-1>config>service>vpls>sap# exit
*A:ALA-1>config>service>vpls#

*A:ALA-2>config>service# vpls 9000 customer 6 create
*A:ALA-2>config>service>vpls# sap 1/1/2:22 create
*A:ALA-2>config>service>vpls>sap$ no shutdown
*A:ALA-2>config>service>vpls>sap# multi-service-site west
*A:ALA-2>config>service>vpls>sap# exit
*A:ALA-2>config>service>vpls#

*A:ALA-3>config>service# vpls 9000 customer 6 create
*A:ALA-3>config>service>vpls# sap 1/1/3:33 create
*A:ALA-3>config>service>vpls>sap$ no shutdown
*A:ALA-3>config>service>vpls>sap# multi-service-site west
*A:ALA-3>config>service>vpls>sap# exit
*A:ALA-3>config>service>vpls#
```

The following example displays a sample configuration of the VPLS SAPs in preparation of a VPLS service between ALA-1, ALA-2, and ALA-3.

```
*A:ALA-1>config>service# info
----------------------------------------
...
        vpls 9000 customer 6 vpn 750 create
            description "Distributed VPLS services."
           def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/2/5:0 create
                description "VPLS SAP"
                multi-service-site "West"
            exit
        exit
```

```
...
-------------------------------------------
*A:ALA-1>config>service#


*A:ALA-2>config>service# info
-------------------------------------------
...
        vpls 9000 customer 6 vpn 750 create
            description "Distributed VPLS services."
           def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/1/2:22 create
                description "VPLS SAP"
                multi-service-site "West"
            exit
        exit
...
-------------------------------------------
*A:ALA-2>config>service#


*A:ALA-3>config>service# info
-----------------------------------------------
...
        vpls 9000 customer 6 vpn 750 create
            description "Distributed VPLS services."
           def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/1/3:33 create
                description "VPLS SAP"
                multi-service-site "West"
            exit
        exit
...
-----------------------------------------------
*A:ALA-3>config>service#
```

Configuring SAP-Specific STP Parameters

When a VPLS has STP enabled, each SAP within the VPLS has STP enabled by default. The operation of STP on each SAP is governed by:

SAP STP Administrative State

The administrative state of STP within a SAP controls how BPDUs are transmitted and handled when received. The allowable states are:

- SAP Admin Up

  The default administrative state is *up* for STP on a SAP. BPDUs are handled in the normal STP manner on a SAP that is administratively up.

- SAP Admin Down

  An administratively down state allows a service provider to prevent a SAP from becoming operationally blocked. BPDUs will not originate out the SAP towards the customer.

  If STP is enabled on VPLS level, but disabled on the SAP, received BPDUs are discarded. Discarding the incoming BPDUs allows STP to continue to operate normally within the VPLS service while ignoring the down SAP. The specified SAP will always be in an operationally forwarding state.

  NOTE: The administratively down state allows a loop to form within the VPLS.

CLI Syntax: `config>service>vpls>sap>stp#`
`[no] shutdown`
Range: `shutdown` or `no shutdown`
Default: `no shutdown` (SAP admin up)

SAP Virtual Port Number

The virtual port number uniquely identifies a SAP within configuration BPDUs. The internal representation of a SAP is unique to a system and has a reference space much bigger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a SAP and identifies it with it's own virtual port number that is unique to every other SAP defined on the VPLS. The virtual port number is assigned at the time that the SAP is added to the VPLS.

Since the order in which SAPs are added to the VPLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance. To achieve consistency after a reboot, the virtual port number can be specified explicitly.

CLI Syntax: `config>service>vpls>sap# stp`
`port-num` *`number`*
Range: 1 — 2047
Default: (automatically generated)
Restore Default: `no port-num`

SAP Priority

SAP priority allows a configurable "tie breaking" parameter to be associated with a SAP. When configuration BPDUs are being received, the configured SAP priority will be used in some circumstances to determine whether a SAP will be designated or blocked. These are the values used for CIST when running MSTP.

In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (0 to 255 also) to create a 16 bit value. In the latest STP standard (802.1D-2004) only the upper 4 bits of the port priority field are used to encode the SAP priority. The remaining 4 bits are used to extend the port ID field into a 12 bit virtual port number field. The virtual port number uniquely references a SAP within the STP instance. See SAP Virtual Port Number on page 390 for details on the virtual port number.

STP computes the actual SAP priority by taking the configured priority value and masking out the lower four bits. The result is the value that is stored in the SAP priority parameter. For example, if a value of 0 was entered, masking out the lower 4 bits would result in a parameter value of 0. If a value of 255 was entered, the result would be 240.

The default value for SAP priority is 128. This parameter can be modified within a range of 0 to 255, 0 being the highest priority. Masking causes the values actually stored and displayed to be 0 to 240, in increments of 16.

CLI Syntax: `config>service>vpls>sap>stp#`
`priority` *`stp-priority`*
Range: 0 to 255 (240 largest value, in increments of 16)
Default: 128
Restore Default: `no priority`

SAP Path Cost

The SAP path cost is used by STP to calculate the path cost to the root bridge. The path cost in BPDUs received on the root port is incremented with the configured path cost for that SAP. When BPDUs are sent out other egress SAPs, the newly calculated root path cost is used. These are the values used for CIST when running MSTP.

STP suggests that the path cost is defined as a function of the link bandwidth. Since SAPs are controlled by complex queuing dynamics, in the 7750 SR the STP path cost is a purely static configuration.

The default value for SAP path cost is 10. This parameter can be modified within a range of 1 to 65535, 1 being the lowest cost.

CLI Syntax: `config>service>vpls>sap>stp#`
`path-cost` *`sap-path-cost`*
Range: 1 to 200000000
Default: 10
Restore Default: `no path-cost`

SAP Edge Port

The SAP `edge-port` command is used to reduce the time it takes a SAP to reach the forwarding state when the SAP is on the edge of the network, and thus has no further STP bridge to handshake with.

The `edge-port` command is used to initialize the internal OPER_EDGE variable. At any time, when OPER_EDGE is false on a SAP, the normal mechanisms are used to transition to the forwarding state (see Forward Delay on page 382). When OPER_EDGE is true, STP assumes that the remote end agrees to transition to the forwarding state without actually receiving a BPDU with an agreement flag set.

The OPER_EDGE variable will dynamically be set to false if the SAP receives BPDUs (the configured edge-port value does not change). The OPER_EDGE variable will dynamically be set to true if auto-edge is enabled and STP concludes there is no bridge behind the SAP.

When STP on the SAP is administratively disabled and re-enabled, the OPER_EDGE is re-initialized to the value configured for edge-port.

Valid values for SAP edge-port are enabled and disabled with disabled being the default.

CLI Syntax: `config>service>vpls>sap>stp#`
`[no] edge-port`
Default: `no edge-port`

SAP Auto Edge

The SAP `edge-port` command is used to instruct STP to dynamically decide whether the SAP is connected to another bridge.

If auto-edge is enabled, and STP concludes there is no bridge behind the SAP, the OPER_EDGE variable will dynamically be set to true. If auto-edge is enabled, and a BPDU is received, the OPER_EDGE variable will dynamically be set to true (see SAP Edge Port on page 391).

Valid values for SAP auto-edge are enabled and disabled with enabled being the default.

CLI Syntax: `config>service>vpls>sap>stp#`
`[no] auto-edge`
Default: `auto-edge`

SAP Link Type

The SAP link-type parameter instructs STP on the maximum number of bridges behind this SAP. If there is only a single bridge, transitioning to forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected via a shared media, their SAPs should all be configured as shared, and timer-based transitions are used.

Valid values for SAP link-type are shared and pt-pt with pt-pt being the default.

CLI Syntax: `config>service>vpls>sap>stp#`
`link-type {pt-pt|shared}`
Default: `link-type pt-pt`
Restore Default: `no link-type`

STP SAP Operational States

The operational state of STP within a SAP controls how BPDUs are transmitted and handled when received. Defined states are:

- Operationally Disabled on page 393
- Operationally Discarding on page 393
- Operationally Learning on page 393
- Operationally Forwarding on page 394

Operationally Disabled

Operationally disabled is the normal operational state for STP on a SAP in a VPLS that has any of the following conditions:

- VPLS state administratively down
- SAP state administratively down
- SAP state operationally down

If the SAP enters the operationally up state with the STP administratively up and the SAP STP state is up, the SAP will transition to the STP SAP discarding state.

When, during normal operation, the router detects a downstream loop behind a SAP or spoke SDP, BPDUs can be received at a very high rate. To recover from this situation, STP will transition the SAP to disabled state for the configured forward-delay duration.

Operationally Discarding

A SAP in the discarding state only receives and sends BPDUs, building the local proper STP state for each SAP while not forwarding actual user traffic. The duration of the discarding state is explained in section Forward Delay on page 382.

Note: in previous versions of the STP standard, the discarding state was called a blocked state.

Operationally Learning

The learning state allows population of the MAC forwarding table before entering the forwarding state. In this state, no user traffic is forwarded.

Operationally Forwarding

Configuration BPDUs are sent out a SAP in the forwarding state. Layer 2 frames received on the SAP are source learned and destination forwarded according to the FIB. Layer 2 frames received on other forwarding interfaces and destined for the SAP are also forwarded.

SAP BPDU Encapsulation State

IEEE 802.1d (referred as dot1d) and Cisco's per VLAN Spanning Tree (PVST) BPDU encapsulations are supported on a per SAP basis. STP is associated with a VPLS service like PVST is per VLAN. The main difference resides in the Ethernet and LLC framing and a type-length-value (TLV) field trailing the BPDU.

Table 14 shows differences between dot1d and PVST Ethernet BPDU encapsulations based on the interface encap-type field:

Table 2: Dot1d and PVST Encapsulation Differences

Field	dot1d encap-type null	dot1d encap-type dot1q	PVST encap-type null	PVST encap-type dot1q
Destination MAC	01:80:c2:00:00:00	01:80:c2:00:00:00	N/A	01:00:0c:cc:cc:cd
Source MAC	Sending Port MAC	Sending Port MAC	N/A	Sending Port MAC
EtherType	N/A	0x81 00	N/A	0x81 00
Dot1p and CFI	N/A	0xe	N/A	0xe
Dot1q	N/A	VPLS SAP ID	N/A	VPLS SAP encap value
Length	LLC Length	LLC Length	N/A	LLC Length
LLC DSAP SSAP	0x4242	0x4242	N/A	0xaaaa (SNAP)
LLC CNTL	0x03	0x03	N/A	0x03
SNAP OUI	N/A	N/A	N/A	00 00 0c (Cisco OUI)
SNAP PID	N/A	N/A	N/A	01 0b
CONFIG	Standard 802.1d	Standard 802.1d	N/A	Standard 802.1d
TLV: Type & Len	N/A	N/A	N/A	58 00 00 00 02
TLV: VLAN	N/A	N/A	N/A	VPLS SAP encap value
Padding	As Required	As Required	N/A	As Required

Each SAP has a Read Only operational state that shows which BPDU encapsulation is currently active on the SAP. The following states apply:

- Dot1d specifies that the switch is currently sending IEEE 802.1d standard BPDUs. The BPDUs will be tagged or non-tagged based on the encapsulation type of the egress interface and the encapsulation value defined in the SAP. A SAP defined on an interface with encapsulation type dot1q will continue in the dot1d BPDU encapsulation state until a PVST encapsulated BPDU is received, in which case the SAP will convert to the PVST encapsulation state. Each received BPDU must be properly IEEE 802.1q tagged if the interface encapsulation type is defined as dot1q. PVST BPDUs will be silently discarded if received when the SAP is on an interface defined with encapsulation type null.
- PVST specifies that the switch is currently sending proprietary encapsulated BPDUs. PVST BPDUs are only supported on Ethernet interfaces with the encapsulation type set to dot1q. The SAP continues in the PVST BPDU encapsulation state until a dot1d encapsulated BPDU is received, in which case the SAP reverts to the dot1d encapsulation state. Each received BPDU must be properly IEEE 802.1q tagged with the encapsulation value defined for the SAP. PVST BPDUs are silently discarded if received when the SAP is on an interface defined with a null encapsulation type.

Dot1d is the initial and only SAP BPDU encapsulation state for SAPs defined on Ethernet interface with encapsulation type set to null.

Each transition between encapsulation types optionally generates an alarm that can be logged and optionally transmitted as an SNMP trap.

Configuring VPLS SAPs with Split Horizon

To configure a VPLS service with a split horizon group, add the split-horizon-group parameter when creating the SAP. Traffic arriving on a SAP within a split horizon group will not be copied to other SAPs in the same split horizon group.

Example: config>service# vpls 800
```
     config>service>vpls# sap 1/1/3:100 split-horizon-group DSL-group1
create
     config>service>vpls>sap$ no shutdown
     config>service>vpls>sap# description "SAP for residential bridging"
     config>service>vpls>sap# exit
     config>service>vpls# sap 1/1/3:200 split-horizon-group DSL-group1
create
     config>service>vpls>sap$ no shutdown
     config>service>vpls>sap# description "SAP for residential bridging"
     config>service>vpls>sap# exit
     config>service>vpls#split-horizon-group DSL-group1
     config>service>vpls>split-horizon-group#description "Split horizon
group for DSL"
     config>service>vpls>split-horizon-group#exit
```

The following example displays the VPLS configuration with split horizon enabled:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
     vpls 800 customer 6001 vpn 700 create
          description "VPLS with split horizon for DSL"
          stp
               shutdown
          exit
          sap 1/1/3:100 split-horizon-group DSL-group1 create
               description "SAP for residential bridging"
          exit
          sap 1/1/3:200 split-horizon-group DSL-group1 create
               description "SAP for residential bridging"
          exit
          split-horizon-group DSL-group1
               description "Split horizon group for DSL"
          exit
          no shutdown
     exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

Configuring MAC Learning Protection

To configure MAC learning protection, configure split horizon, MAC protection, and SAP parameters.

Example: `config>service# vpls 800`

```
config>service>vpls# split-horizon-group "DSL-group1"
config>service>vpls>split-horizon-group# restrict-protected-src
config>service>vpls>split-horizon-group# restrict-unprotected-dst
config>service>vpls>split-horizon-group# exit
config>service>vpls# mac-protect
config>service>vpls>mac-protect# mac FF:FF:FF:FF:FF:FF
config>service>vpls>mac-protect# exit
config>service>vpls# sap 1/1/3:100
config>service>vpls>sap# restrict-protected-src
```

The following example displays the VPLS configuration with split horizon enabled:

```
A:ALA-48>config>service>vpls# info
----------------------------------------------
            description "IMA VPLS"
            split-horizon-group "DSL-group1" create
                restrict-protected-src
                restrict-unprotected-dst
            exit
            mac-protect
                 mac ff:ff:ff:ff:ff:ff
            exit
            sap 1/1/9:0 create
                ingress
                    scheduler-policy "SLA1"
                    qos 100 shared-queuing
                exit
                egress
                    scheduler-policy "SLA1"
                    filter ip 10
                exit
                restrict-protected-src
                arp-reply-agent
                host-connectivity-verify source-ip 143.144.145.1
            exit
...
----------------------------------------------
A:ALA-48>config>service>vpls#
```

Applying an Egress Multicast Group to a VPLS Service SAP

Use the following CLI syntax to apply an egress multicast group to a VPLS service SAP:

CLI Syntax: config>service>vpls *service-id* [customer *customer-id*] [vpn *vpn-id*] [mvpls]
sap *sap-id* [split-horizon-group *group-name*]
egress
multicast-group *group-name*

Example: config>service# vpls 800
config>service>vpls# sap 1/1/3:100
config>service>vpls>sap# egress
config>service>vpls>sap>egress# multicast-group "vpls-emg-1"
config>service>vpls>sap>egress# exit
config>service>vpls>sap# exit

The following example displays the VPLS configuration with egress multicast group:

```
A:ALA-48>config>service>vpls# info
----------------------------------------------
            description "VPLS with split horizon for DSL"
            split-horizon-group "DSL-group1" create
                description "Split horizon group for DSL"
            exit
            stp
                shutdown
            exit
            sap 1/1/4:200 split-horizon-group "DSL-group1" create
                description "SAP for residential bridging"
            exit
            sap 1/1/3:100 split-horizon-group "DSL-group1" create
                description "SAP for residential bridging"
                egress
                   multicast-group "vpls-emg-1"
            exit
            no shutdown
----------------------------------------------
A:ALA-48>config>service>vpls#
```

Configuring SAP Subscriber Management Parameters

Use the following CLI syntax to configure subscriber management parameters on a VPLS service SAP. The policies and profiles that are referenced in the the `def-sla-profile`, `def-sub-profile`, `non-sub-traffic` , and `sub-ident-policy` commands must already be configured in the **config>subscr-mgmt** context.

CLI Syntax:
```
config>service>vpls service-id
    sap sap-id [split-horizon-group group-name]
        sub-sla-mgmt
            def-sla-profile default-sla-profile-name
            def-sub-profile default-subscriber-profile-name
            mac-da-hashing
            multi-sub-sap [number-of-sub]
            no shutdown
            single-sub-parameters
                non-sub-traffic sub-profile sub-profile-name sla-
                    profile sla-profile-name [subscriber sub-ident-
                    string]
                profiled-traffic-only
            sub-ident-policy sub-ident-policy-name
```

Example:
```
config>service# vpls 90001
    config>service>vpls# sap 1/2/2:0
    config>service>vpls>sap# sub-sla-mgmt
    config>service>vpls>sap>sub-sla-mgmt# def-sla-profile sla-profile1
    config>service>vpls>sap>sub-sla-mgmt# sub-ident-policy "SubIdent1"
    config>service>vpls>sap>sub-sla-mgmt# exit
    config>service>vpls>sap# exit
    config>service>vpls#
```

The following example displays the subscriber management configuration:

```
A:ALA-48>config>service>vpls#
------------------------------------------
            description "Local VPLS"
            stp
                shutdown
            exit
            sap 1/2/2:0 create
                description "SAP for local service"
                sub-sla-mgmt
                    def-sla-profile "sla-profile1"
                    sub-ident-policy "SubIdent1"
                exit
            exit
            sap 1/1/5:0 create
                description "SAP for local service"
            exit
            no shutdown
------------------------------------------
A:ALA-48>config>service>vpls#
```

Configuring SDP Bindings

VPLS provides scaling and operational advantages. A hierarchical configuration eliminates the need for a full mesh of VCs between participating devices. Hierarchy is achieved by enhancing the base VPLS core mesh of VCs with access VCs (spoke) to form two tiers. Spoke SDPs are generally created between Layer 2 switches and placed at the Multi-Tenant Unit (MTU). The PE routers are placed at the service provider's Point of Presence (POP). Signaling and replication overhead on all devices is considerably reduced.

A spoke SDP is treated like the equivalent of a traditional bridge port where flooded traffic received on the spoke SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received (unless a split horizon group was defined on the spoke SDP, see section Configuring VPLS spoke SDPs with Split Horizon on page 411).

A spoke SDP connects a VPLS service between two sites and, in its simplest form, could be a single tunnel LSP. A set of ingress and egress VC labels are exchanged for each VPLS service instance to be transported over this LSP. The PE routers at each end treat this as a virtual spoke connection for the VPLS service in the same way as the PE-MTU connections. This architecture minimizes the signaling overhead and avoids a full mesh of VCs and LSPs between the two metro networks.

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

A VC-ID can be specified with the SDP-ID. The VC-ID is used instead of a label to identify a virtual circuit. The VC-ID is significant between peer SRs on the same hierarchical level. The value of a VC-ID is conceptually independent from the value of the label or any other datalink specific information of the VC.

Figure 43 displays an example of a distributed VPLS service configuration of spoke and mesh SDPs (uni-directional tunnels) between SR-Series routers and MTUs.

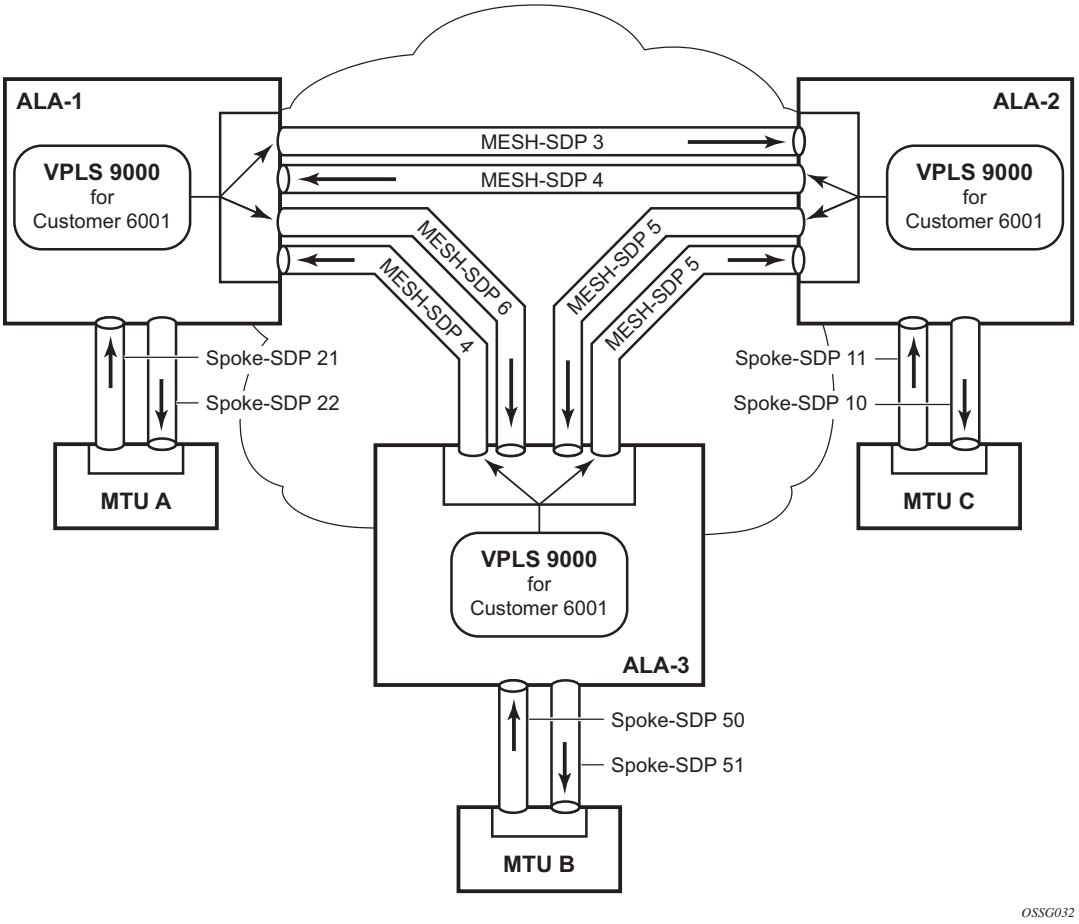

Figure 1: SDPs — Uni-Directional Tunnels

Use the following CLI syntax to create a mesh or spoke SDP bindings with a distributed VPLS service. SDPs must be configured prior to binding. Refer to Configuring an SDP on page 68 for information about creating SDPs.

Use the following CLI syntax to configure mesh SDP bindings:

CLI Syntax:

```
config>service# vpls service-id
 mesh-sdp sdp-id[:vc-id] [vc-type {ether|vlan}]
    egress
       filter {ip ip-filter-id|mac mac-filter-id}
       mfib-allowed-mda-destinations
             mda mda-id
       vc-label egress-vc-label
    ingress
       filter {ip ip-filter-id|mac mac-filter-id}
       vc-label ingress-vc-label
    no shutdown
    static-mac ieee-address
    vlan-vc-tag 0..4094
```

Use the following CLI syntax to configure spoke SDP bindings:

CLI Syntax:

```
config>service# vpls service-id
 spoke-sdp sdp-id:vc-id [vc-type {ether|vlan}] [split-hori-
    zon-group group-name]
    egress
       filter {ip ip-filter-id|mac mac-filter-id}
       vc-label egress-vc-label
    ingress
       filter {ip ip-filter-id|mac mac-filter-id}
       vc-label ingress-vc-label
    limit-mac-move[non-blockable]
    vlan-vc-tag 0..4094
    no shutdown
    static-mac ieee-address
    stp
       path-cost stp-path-cost
       priority stp-priority
       no shutdown
    vlan-vc-tag [0..4094]
```

The following example displays the command usage to configure the spoke and mesh SDPs displayed in Figure 43 on page 401.

Example:

```
*A:ALA-1>config>service# vpls 9000
*A:ALA-1>config>service>vpls# mesh-sdp 5 create
*A:ALA-1>config>service>vpls> mesh-sdp# no shutdown
*A:ALA-1>config>service>vpls> mesh-sdp# exit
*A:ALA-1>config>service>vpls# mesh-sdp 7 create
*A:ALA-1>config>service>vpls> mesh-sdp# no shutdown
*A:ALA-1>config>service>vpls> mesh-sdp# exit
*A:ALA-1>config>service>vpls# spoke-sdp 2:22 create
*A:ALA-1>config>service>vpls> spoke-sdp# no shutdown
*A:ALA-1>config>service>vpls> spoke-sdp# exit
*A:ALA-1>config>service>vpls#

*A:ALA-2>config>service# vpls 9000
*A:ALA-2>config>service>vpls# mesh-sdp 5 create
*A:ALA-2>config>service>vpls> mesh-sdp# no shutdown
*A:ALA-2>config>service>vpls> mesh-sdp# exit
*A:ALA-2>config>service>vpls# mesh-sdp 7 create
*A:ALA-2>config>service>vpls> mesh-sdp# no shutdown
*A:ALA-2>config>service>vpls> mesh-sdp# exit
*A:ALA-2>config>service>vpls# spoke-sdp 2:22 create
*A:ALA-2>config>service>vpls> spoke-sdp# no shutdown
*A:ALA-2>config>service>vpls> spoke-sdp# exit
*A:ALA-2>config>service>vpls#
```

```
*A:ALA-3>config>service# vpls 9000
*A:ALA-3>config>service>vpls# mesh-sdp 5 create
*A:ALA-3>config>service>vpls> mesh-sdp# no shutdown
*A:ALA-3>config>service>vpls> mesh-sdp# exit
*A:ALA-3>config>service>vpls# mesh-sdp 7 create
*A:ALA-3>config>service>vpls> mesh-sdp# no shutdown
*A:ALA-3>config>service>vpls> mesh-sdp# exit
*A:ALA-3>config>service>vpls# spoke-sdp 2:22 create
*A:ALA-3>config>service>vpls> spoke-sdp# no shutdown
*A:ALA-3>config>service>vpls> spoke-sdp# exit
*A:ALA-3>config>service>vpls#
```

The following displays the SDP binding configurations for ALA-1, ALA-2, and ALA-3 for VPLS service ID 9000 for customer 6:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
            def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/2/5:0 create
            exit
            spoke-sdp 2:22 create
            exit
            mesh-sdp 5:750 create
            exit
            mesh-sdp 7:750 create
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-1>config>service#


*A:ALA-2>config>service# info
----------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
            def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/1/2:22 create
            exit
            spoke-sdp 2:22 create
            exit
            mesh-sdp 5:750 create
            exit
            mesh-sdp 7:750 create
            exit
            no shutdown
        exit
----------------------------------------------
```

```
*A:ALA-3>config>service# info
----------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
            def-mesh-vc-id 750
            stp
                shutdown
            exit
            sap 1/1/3:33 create
            exit
            spoke-sdp 2:22 create
            exit
            mesh-sdp 5:750 create
            exit
            mesh-sdp 7:750 create
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-3>config>service#
```

Configuring Spoke SDP Specific STP Parameters

When a VPLS has STP enabled, each spoke SDP within the VPLS has STP enabled by default. The operation of STP on each spoke SDP is governed by:

Spoke SDP STP Administrative State

The administrative state of STP within a spoke SDP controls how BPDUs are transmitted and handled when received. The allowable states are:

- Spoke SDP Admin Up

  The default administrative state is *up* for STP on a spoke SDP. BPDUs are handled in the normal STP manner on a spoke SDP that is administratively up.

- Spoke SDP Admin Down

  An administratively down state allows a service provider to prevent a spoke SDP from becoming operationally blocked. BPDUs will not originate out the spoke SDP towards the customer.

  If STP is enabled on VPLS level, but disabled on the spoke SDP, received BPDUs are discarded. Discarding the incoming BPDUs allows STP to continue to operate normally within the VPLS service while ignoring the down spoke SDP. The specified spoke SDP will always be in an operationally forwarding state.

  NOTE: The administratively down state allows a loop to form within the VPLS.

CLI Syntax: `config>service>vpls>spoke-sdp>stp#`
`[no] shutdown`
Range: `shutdown` or `no shutdown`
Default: `no shutdown` (spoke SDP admin up)

Spoke SDP Virtual Port Number

The virtual port number uniquely identifies a spoke SDP within configuration BPDUs. The internal representation of a spoke SDP is unique to a system and has a reference space much

bigger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a spoke SDP and identifies it with it's own virtual port number that is unique to every other spoke SDP defined on the VPLS. The virtual port number is assigned at the time that the spoke SDP is added to the VPLS.

Since the order in which spoke SDPs are added to the VPLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance. To achieve consistency after a reboot, the virtual port number can be specified explicitly.

CLI Syntax: `config>service>vpls>spoke-sdp# stp`
`port-num` *`number`*
Range: 1 — 2047
Default: (automatically generated)
Restore Default: `no port-num`

Spoke SDP Priority

Spoke SDP priority allows a configurable tie breaking parameter to be associated with a spoke SDP. When configuration BPDUs are being received, the configured spoke SDP priority will be used in some circumstances to determine whether a spoke SDP will be designated or blocked.

In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (0 to 255 also) to create a 16 bit value. In the latest STP standard (802.1D-2004) only the upper 4 bits of the port priority field are used to encode the spoke SDP priority. The remaining 4 bits are used to extend the port ID field into a 12 bit virtual port number field. The virtual port number uniquely references a spoke SDP within the STP instance. See Spoke SDP Virtual Port Number on page 405 for details on the virtual port number.

STP computes the actual spoke SDP priority by taking the configured priority value and masking out the lower four bits. The result is the value that is stored in the spoke SDP priority parameter. For instance, if a value of 0 was entered, masking out the lower 4 bits would result in a parameter value of 0. If a value of 255 was entered, the result would be 240.

The default value for spoke SDP priority is 128. This parameter can be modified within a range of 0 to 255, 0 being the highest priority. Masking causes the values actually stored and displayed to be 0 to 240, in increments of 16.

CLI Syntax: `config>service>vpls>spoke-sdp>stp#`
`priority` *`stp-priority`*
Range: 0 to 255 (240 largest value, in increments of 16)
Default: 128
Restore Default: `no priority`

Spoke SDP Path Cost

The spoke SDP path cost is used by STP to calculate the path cost to the root bridge. The path cost in BPDUs received on the root port is incremented with the configured path cost for that spoke SDP. When BPDUs are sent out other egress spoke SDPs, the newly calculated root path cost is used.

STP suggests that the path cost is defined as a function of the link bandwidth. Since spoke SDPs are controlled by complex queuing dynamics, the STP path cost is a purely static configuration.

The default value for spoke SDP path cost is 10. This parameter can be modified within a range of 1 to 200000000 (1 is the lowest cost).

CLI Syntax: `config>service>vpls>spoke-sdp>stp#`
`path-cost` *`stp-path-cost`*
Range: 1 to 200000000
Default: 10
Restore Default: `no path-cost`

Spoke SDP Edge Port

The spoke SDP `edge-port` command is used to reduce the time it takes a spoke SDP to reach the forwarding state when the spoke SDP is on the edge of the network, and thus has no further STP bridge to handshake with.

The `edge-port` command is used to initialize the internal OPER_EDGE variable. At any time, when OPER_EDGE is false on a spoke SDP, the normal mechanisms are used to transition to the forwarding state (see Forward Delay on page 382). When OPER_EDGE is true, STP assumes that the remote end agrees to transition to the forwarding state without actually receiving a BPDU with an agreement flag set.

The OPER_EDGE variable will dynamically be set to false if the spoke SDP receives BPDUs (the configured edge-port value does not change). The OPER_EDGE variable will dynamically be set to true if auto-edge is enabled and STP concludes there is no bridge behind the spoke SDP.

When STP on the spoke SDP is administratively disabled and re-enabled, the OPER_EDGE is re-initialized to the spoke SDP configured for edge-port.

Valid values for spoke SDP edge-port are enabled and disabled with disabled being the default.

CLI Syntax: `config>service>vpls>spoke-sdp>stp#`
`[no] edge-port`
Default: `no edge-port`

Spoke SDP Auto Edge

The spoke SDP `edge-port` command is used to instruct STP to dynamically decide whether the spoke SDP is connected to another bridge.

If auto-edge is enabled, and STP concludes there is no bridge behind the spoke SDP, the OPER_EDGE variable will dynamically be set to true. If auto-edge is enabled, and a BPDU is received, the OPER_EDGE variable will dynamically be set to true (see Spoke SDP Edge Port on page 407).

Valid values for spoke SDP auto-edge are enabled and disabled with enabled being the default.

CLI Syntax: `config>service>vpls>spoke-sdp>stp#`
`[no] auto-edge`
Default: `auto-edge`

Spoke SDP Link Type

The spoke SDP `link-type` command instructs STP on the maximum number of bridges behind this spoke SDP. If there is only a single bridge, transitioning to forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected via a shared media, their spoke SDPs should all be configured as shared, and timer-based transitions are used.

Valid values for spoke SDP link-type are shared and pt-pt with pt-pt being the default.

CLI Syntax: `config>service>vpls>spoke-sdp>stp#`
`link-type {pt-pt|shared}`
Default: `link-type pt-pt`
Restore Default: `no link-type`

Spoke SDP STP Operational States

The operational state of STP within a spoke SDP controls how BPDUs are transmitted and handled when received. Defined states are:

- Operationally Disabled on page 409
- Operationally Discarding on page 409
- Operationally Learning on page 409
- Operationally Forwarding on page 410

Operationally Disabled

Operationally disabled is the normal operational state for STP on a spoke SDP in a VPLS that has any of the following conditions:

- VPLS state administratively down
- Spoke SDP state administratively down
- Spoke SDP state operationally down

If the spoke SDP enters the operationally up state with the STP administratively up and the spoke SDP STP state is up, the spoke SDP will transition to the STP spoke SDP discarding state.

When, during normal operation, the router detects a downstream loop behind a spoke SDP, BPDUs can be received at a very high rate. To recover from this situation, STP will transition the spoke SDP to a disabled state for the configured forward-delay duration.

Operationally Discarding

A spoke SDP in the discarding state only receives and sends BPDUs, building the local proper STP state for each spoke SDP while not forwarding actual user traffic. The duration of the discarding state is explained in section Forward Delay on page 382.

Note: in previous versions of the STP standard, the discarding state was called a blocked state.

Operationally Learning

The learning state allows population of the MAC forwarding table before entering the forwarding state. In this state no user traffic is forwarded.

Operationally Forwarding

Configuration BPDUs are sent out a spoke SDP in the forwarding state. Layer 2 frames received on the spoke SDP are source learned and destination forwarded according to the FIB. Layer 2 frames received on other forwarding interfaces and destined for the spoke SDP are also forwarded.

Spoke SDP BPDU Encapsulation States

IEEE 802.1D (referred as dot1d) and Cisco's per VLAN Spanning Tree (PVST) BPDU encapsulations are supported on a per spoke SDP basis. STP is associated with a VPLS service like PVST is per VLAN. The main difference resides in the Ethernet and LLC framing and a type-length-value (TLV) field trailing the BPDU.

Table 14 shows differences between dot1D and PVST Ethernet BPDU encapsulations based on the interface encap-type field:

Table 3: Dot1d and PVST Encapsulation Differences

Field	dot1d encap-type null	dot1d encap-type dot1q	PVST encap-type null	PVST encap-type dot1q
Destination MAC	01:80:c2:00:00:00	01:80:c2:00:00:00	N/A	01:00:0c:cc:cc:cd
Source MAC	Sending Port MAC	Sending Port MAC	N/A	Sending Port MAC
EtherType	N/A	0x81 00	N/A	0x81 00
Dot1p and CFI	N/A	0xe	N/A	0xe
Dot1q	N/A	VPLS spoke SDP ID	N/A	VPLS spoke SDP encap value
Length	LLC Length	LLC Length	N/A	LLC Length
LLC DSAP SSAP	0x4242	0x4242	N/A	0xaaaa (SNAP)
LLC CNTL	0x03	0x03	N/A	0x03
SNAP OUI	N/A	N/A	N/A	00 00 0c (Cisco OUI)
SNAP PID	N/A	N/A	N/A	01 0b
CONFIG or TCN BPDU	Standard 802.1d	Standard 802.1d	N/A	Standard 802.1d
TLV: Type & Len	N/A	N/A	N/A	58 00 00 00 02
TLV: VLAN	N/A	N/A	N/A	VPLS spoke SDP encap value
Padding	As Required	As Required	N/A	As Required

Each spoke SDP has a Read Only operational state that shows which BPDU encapsulation is currently active on the spoke SDP. The following states apply:

- **Dot1d** specifies that the switch is currently sending IEEE 802.1D standard BPDUs. The BPDUs will be tagged or non-tagged based on the encapsulation type of the egress interface and the encapsulation value defined in the spoke SDP. A spoke SDP defined on an interface with encapsulation type dot1q will continue in the dot1d BPDU encapsulation state until a PVST encapsulated BPDU is received, after which the spoke SDP will convert to the PVST encapsulation state. Each received BPDU must be properly IEEE 802.1q tagged if the interface encapsulation type is defined to dot1q.
- **PVST** specifies that the switch is currently sending proprietary encapsulated BPDUs. PVST BPDUs are only supported on Ethernet interfaces with the encapsulation type set to dot1q. The spoke SDP continues in the PVST BPDU encapsulation state until a dot1d encapsulated BPDU is received, in which case the spoke SDP reverts to the dot1d encapsulation state. Each received BPDU must be properly IEEE 802.1q tagged with the encapsulation value defined for the spoke SDP.

Dot1d is the initial and only spoke SDP BPDU encapsulation state for spoke SDPs defined on Ethernet interface with encapsulation type set to null.

Each transition between encapsulation types optionally generates an alarm that can be logged and optionally transmitted as an SNMP trap.

Configuring VPLS spoke SDPs with Split Horizon

To configure spoke SDPs with a split horizon group, add the split-horizon-group parameter when creating the spoke SDP. Traffic arriving on a SAP or spoke SDP within a split horizon group will not be copied to other SAPs or spoke SDPs in the same split horizon group.

Example: `config>service# vpls 800`

```
     config>service>vpls# spoke-sdp 51:15 split-horizon-group DSL-group1
     config>service>vpls>spoke-sdp $ no shutdown
     config>service>vpls>spoke-sdp # exit
     config>service>vpls#split-horizon-group DSL-group1
     config>service>vpls>split-horizon-group# description "Split horizon
group for DSL"
     config>service>vpls>split-horizon-group#exit
```

The following example displays the VPLS configuration with split horizon enabled:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
vpls 800 customer 6001 vpn 700 create
     description "VPLS with split horizon for DSL"
     stp
          shutdown
     exit
```

```
        spoke-sdp 51:15 split-horizon-group DSL-group1 create
        exit
        split-horizon-group DSL-group1
            description "Split horizon group for DSL"
        exit
        no shutdown
    exit
    ...
    ----------------------------------------------
    *A:ALA-1>config>service#
```

Configuring VPLS Redundancy

This section discusses the following service management tasks:

Creating a Management VPLS for SAP Protection

This section provides a brief overview of the tasks that must be performed to configure a management VPLS for SAP protection and provides the CLI commands, see Figure 44. The tasks below should be performed on both 7750 SR nodes providing the protected VPLS service.

Before configuring a management VPLS, first read VPLS Redundancy on page 350 for an introduction to the concept of management VPLS and SAP redundancy.

1. Create an SDP to the peer node.
2. Create a management VPLS.
3. Define a SAP in the m-vpls on the port towards the MTU. Note that the port must be dot1q or qinq tagged. The SAP corresponds to the (stacked) VLAN on the MTU in which STP is active.

   Optionally modify STP parameters for load balancing (see Configuring Load Balancing with Management VPLS on page 419).
4. Create a mesh SDP in the m-vpls using the SDP defined in step 1. Ensure that this mesh SDP runs over a protected LSP (see note below).
5. Enable the management VPLS service and verify that it is operationally up.
6. Create a list of VLANs on the port that are to be managed by this management VPLS.
7. Create one or more user VPLS services with SAPs on VLANs in the range defined by Step 6.

Note: The mesh SDP should be protected by e.g., a backup LSP or Fast Reroute. If the mesh SDP were to go down, STP on both nodes would go to “forwarding” state and a loop would occur.

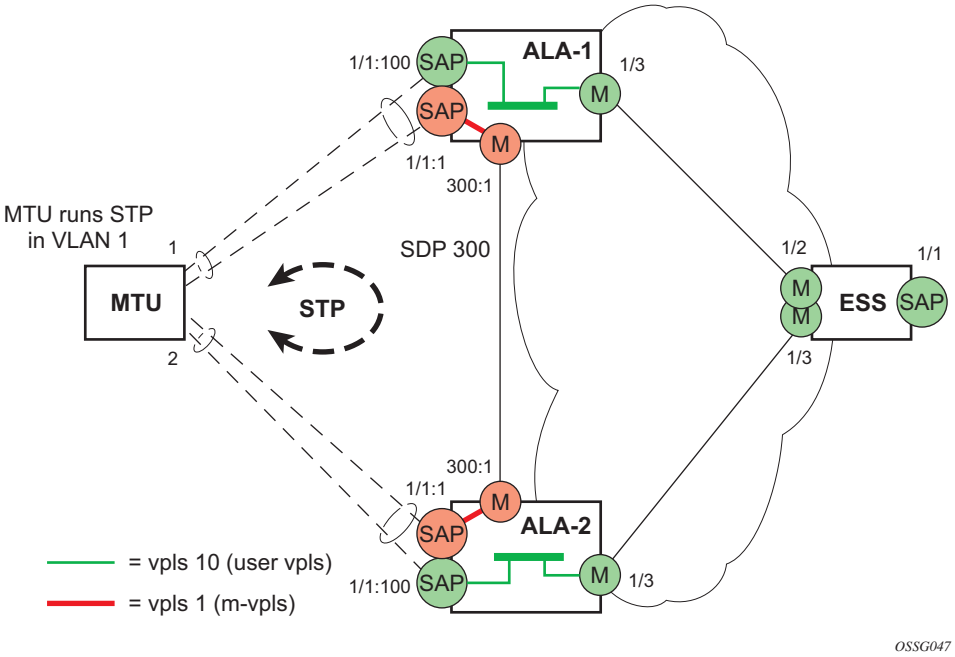

Figure 2: Example Configuration for Protected VPLS SAP

Use the following CLI syntax to create a management VPLS:

CLI Syntax:
```
config>service# sdp sdp-id mpls create
    far-end ip-address
    lsp lsp-name
    no shutdown
```

CLI Syntax:
```
vpls service-id customer customer-id [m-vpls] create
    description description-string
    sap sap-id create
       managed-vlan-list
          range vlan-range
    mesh-sdp sdp-id:vc-id create
    stp
    no shutdown
```

Example:
```
config>service#  sdp 300 mpls create
config>service>sdp# far-end 10.0.0.20
config>service>sdp# lsp "toALA-A2"
config>service>sdp# no shutdown
config>service>sdp# exit
config>service>vpls# mesh-sdp 300:1 create
config>service>vpls>mesh-sdp# exit
config>service>vpls# no shutdown
```

```
config>service# vpls 1 customer 1 m-vpls create
config>service>vpls# sap 1/1/1:1 create
config>service>vpls>sap# managed-vlan-list
config>service>vpls>vlan-list# range 100-1000
config>service>vpls>vlan-list# exit
config>service>vpls>sap# exit
config>service>vpls# no shutdown
config>service>vpls# exit
```

The following example displays the VPLS configuration:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
       sdp 300 mpls create
            far-end 10.0.0.20
            lsp "toALA-A2"
            no shutdown
       exit
       vpls 1 customer 1 m-vpls create
            sap 1/1/1:1 create
                managed-vlan-list
                    range 100-1000
                exit
            exit
            mesh-sdp 300:1 create
            exit
            stp
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

Creating a Management VPLS for Spoke SDP Protection

This section provides a brief overview of the tasks that must be performed to configure a management VPLS for spoke SDP protection and provides the CLI commands, see Figure 45. The tasks below should be performed on all four 7750 SR nodes providing the protected VPLS service.

Before configuring a management VPLS, please first read Configuring a VPLS SAP on page 386 for an introduction to the concept of management VPLS and spoke SDP redundancy.

1. Create an SDP to the local peer 7750 SR node (node ALA-A2 in the example below).
2. Create an SDP to the remote peer 7750 SR node (node ALA-B1 in the example below).
3. Create a management VPLS.
4. Create a mesh SDP in the m-vpls using the SDP defined in Step 1. Ensure that this mesh SDP runs over a protected LSP (see note below).
5. Enable the management VPLS service en verify that it is operationally up.
6. Create a spoke SDP inthe m-vpls using the SDP defined in Step 2.

   Optionally modify STP parameters for load balancing (see Configuring Load Balancing with Management VPLS on page 419).
7. Create one or more user VPLS services with spoke SDPs on the tunnel SDP defined by Step 2.

As long as the user spoke SDPs created in step 8 are in this same tunnel SDP with the management spoke SDP created in step 6, the management VPLS will protect them.

The mesh SDP should be protected by, for example, a backup LSP or Fast Reroute. If the mesh SDP were to go down, STP on both nodes would go to “forwarding” state and a loop would occur.

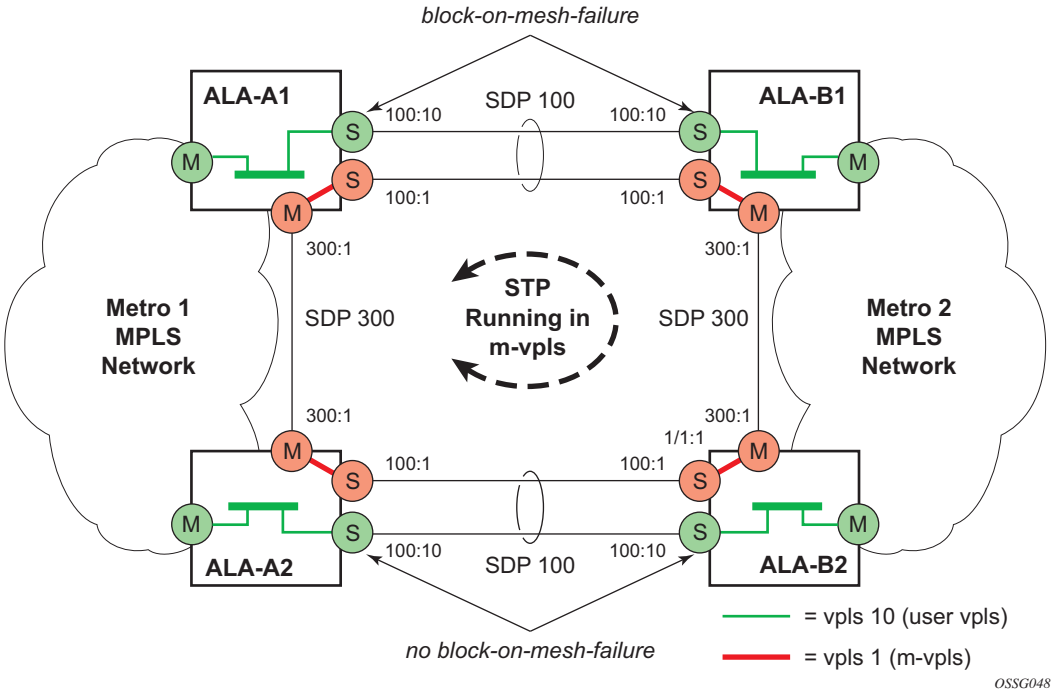

Figure 3: Example Configuration for Protected VPLS Spoke SDP

Use the following CLI syntax to create a management VPLS for spoke SDP protection:

CLI Syntax:
```
config>service# sdp sdp-id mpls create
 far-end ip-address
 lsp lsp-name
 no shutdown
```

CLI Syntax:
```
vpls service-id customer customer-id [m-vpls] create
 description description-string
 mesh-sdp sdp-id:vc-id create
 spoke-sdp sdp-id:vc-id create
 stp
 no shutdown
```

Example:
```
config>service#  sdp 100 mpls create
config>service>sdp# far-end 10.0.0.30
config>service>sdp# lsp "toALA-B1"
config>service>sdp# no shutdown
config>service>sdp# exit
config>service#  sdp 300 mpls create
config>service>sdp# far-end 10.0.0.20
config>service>sdp# lsp "toALA-A2"
config>service>sdp# no shutdown
```

```
config>service>sdp# exit
config>service# vpls 1 customer 1 m-vpls create
config>service>vpls# mesh-sdp 300:1 create
config>service>vpls>mesh-sdp# exit
config>service>vpls# no shutdown
config>service>vpls# spoke-sdp 100:1 create
config>service>vpls>spoke-sdp# exit
config>service>vpls# exit
```

The following example displays the VPLS configuration:

```
*A:ALA-A1>config>service# info
----------------------------------------------
...
       sdp 100 mpls create
            far-end 10.0.0.30
            lsp "toALA-B1"
            no shutdown
       exit
       sdp 300 mpls create
            far-end 10.0.0.20
            lsp "toALA-A2"
            no shutdown
       exit
       vpls 101 customer 1 m-vpls create
            spoke-sdp 100:1 create
            exit
            mesh-sdp 300:1 create
            exit
            stp
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-A1>config>service#
```

Configuring Load Balancing with Management VPLS

With the concept of management VPLS, it is possible to load balance the user VPLS services across the two protecting nodes. This is done by creating two management VPLS instances, where both instances have different active spokes (by changing the STP path-cost). When different user VPLS services are associated with either the two management VPLS services, the traffic will be split across the two spokes.

Load balancing can be achieved in both the SAP protection and spoke SDP protection scenarios.

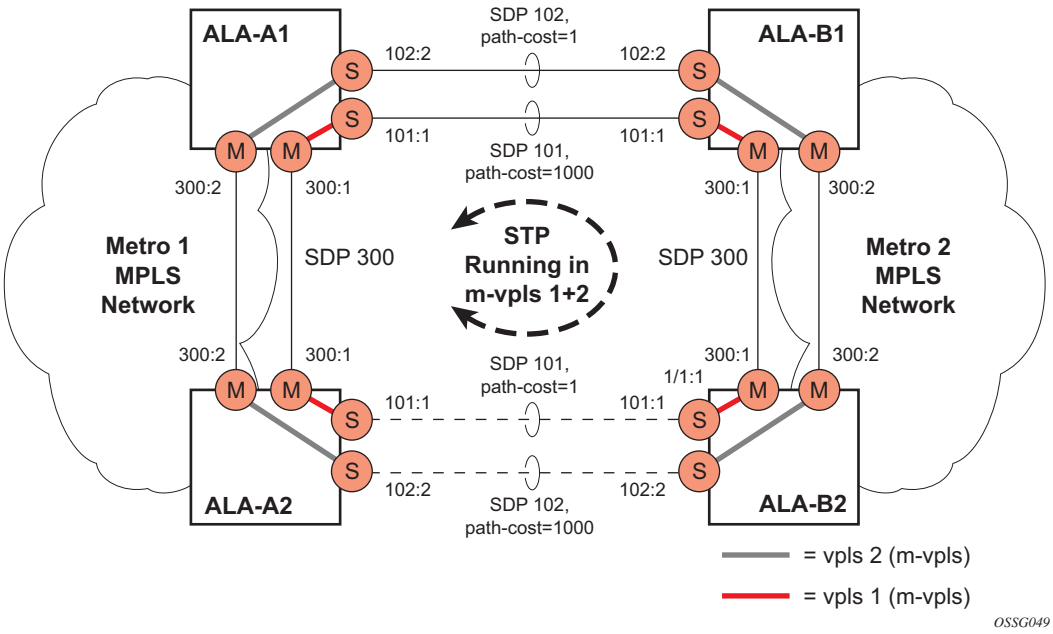

Figure 4: Example Configuration for Loadbalancing Across Two Protected VPLS Spoke SDPs

Use the following CLI syntax to create a load balancing across two management VPLS instances:

CLI Syntax:
```
config>service# sdp sdp-id mpls create
    far-end ip-address
    lsp lsp-name
    no shutdown
```

CLI Syntax:
```
vpls service-id customer customer-id [m-vpls] create
    description description-string
    mesh-sdp sdp-id:vc-id create
    spoke-sdp sdp-id:vc-id create
        stp
```

```
            path-cost
        stp
        no shutdown
```

Example:

```
config>service#  sdp 101 mpls create
config>service>sdp# far-end 10.0.0.30
config>service>sdp# lsp "1toALA-B1"
config>service>sdp# no shutdown
config>service>sdp# exit
config>service#  sdp 102 mpls create
config>service>sdp# far-end 10.0.0.20
config>service>sdp# lsp "2toALA-B1"
config>service>sdp# no shutdown
config>service>sdp# exit
config>service# vpls 1 customer 1 m-vpls create
config>service>vpls# mesh-sdp 300:1 create
config>service>vpls>mesh-sdp# exit
config>service>vpls# no shutdown
config>service>vpls# spoke-sdp 101:1 create
config>service>vpls>spoke-sdp# stp
config>service>vpls>spoke-sdp>stp# path-cost 1
config>service>vpls>spoke-sdp>stp# exit
config>service>vpls>spoke-sdp# exit
config>service>vpls# exit
config>service# vpls 2 customer 1 m-vpls create
config>service>vpls# mesh-sdp 300:2 create
config>service>vpls>mesh-sdp# exit
config>service>vpls# no shutdown
config>service>vpls# spoke-sdp 102:1 create
config>service>vpls>spoke-sdp# stp
config>service>vpls>spoke-sdp>stp# path-cost 1000
config>service>vpls>spoke-sdp>stp# exit
config>service>vpls>spoke-sdp# exit
config>service>vpls# exit
```

Note: the STP path costs in each peer 7750 SR node should be reversed (Figure 46).

The following example displays the VPLS configuration on ALA-A1 (top left, IP address 10.0.0.10):

```
*A:ALA-A1>config>service# info
----------------------------------------------
...
      sdp 101 mpls create
            far-end 10.0.0.30
            lsp "1toALA-B1"
            no shutdown
      exit
      sdp 102 mpls create
            far-end 10.0.0.30
            lsp "2toALA-B1"
            no shutdown
      exit
 ...
     vpls 101 customer 1 m-vpls create
            spoke-sdp 101:1 create
               stp
                  path-cost 1
               exit
            exit
            mesh-sdp 300:1 create
            exit
            stp
            exit
           no shutdown
      exit
      vpls 102 customer 1 m-vpls create
            spoke-sdp 102:2 create
               stp
                  path-cost 1000
               exit
            exit
            mesh-sdp 300:2 create
            exit
            stp
            exit
            no shutdown
      exit
...
----------------------------------------------
*A:ALA-A1>config>service#
```

The following example displays the VPLS configuration on ALA-A2 (bottom left, IP address 10.0.0.20):

```
*A:ALA-A2>config>service# info
----------------------------------------------
...
      sdp 101 mpls create
            far-end 10.0.0.40
            lsp "1toALA-B2"
            no shutdown
      exit
      sdp 102 mpls create
            far-end 10.0.0.40
            lsp "2toALA-B2"
            no shutdown
      exit
 ...
     vpls 101 customer 1 m-vpls create
            spoke-sdp 101:1 create
               stp
                  path-cost 1000
               exit
            exit
            mesh-sdp 300:1 create
            exit
            stp
            exit
           no shutdown
      exit
      vpls 102 customer 1 m-vpls create
            spoke-sdp 102:2 create
               stp
                  path-cost 1
               exit
            exit
            mesh-sdp 300:2 create
            exit
            stp
            exit
            no shutdown
      exit
...
----------------------------------------------
*A:ALA-A2>config>service#
```

The following example displays the VPLS configuration on ALA-A3 (top right, IP address 10.0.0.30):

```
*A:ALA-A1>config>service# info
----------------------------------------------
...
      sdp 101 mpls create
            far-end 10.0.0.10
            lsp "1toALA-A1"
            no shutdown
      exit
      sdp 102 mpls create
            far-end 10.0.0.10
            lsp "2toALA-A1"
            no shutdown
      exit
 ...
     vpls 101 customer 1 m-vpls create
            spoke-sdp 101:1 create
               stp
                  path-cost 1
               exit
            exit
            mesh-sdp 300:1 create
            exit
            stp
            exit
            no shutdown
      exit
      vpls 102 customer 1 m-vpls create
            spoke-sdp 102:2 create
               stp
                  path-cost 1000
               exit
            exit
            mesh-sdp 300:2 create
            exit
            stp
            exit
            no shutdown
      exit
...
----------------------------------------------
*A:ALA-A1>config>service#
```

The following example displays the VPLS configuration on ALA-A4 (bottom right, IP address 10.0.0.40):

```
*A:ALA-A2>config>service# info
----------------------------------------------
...
      sdp 101 mpls create
            far-end 10.0.0.20
            lsp "1toALA-B2"
            no shutdown
      exit
      sdp 102 mpls create
            far-end 10.0.0.20
            lsp "2toALA-B2"
            no shutdown
      exit
 ...
     vpls 101 customer 1 m-vpls create
            spoke-sdp 101:1 create
               stp
                  path-cost 1000
               exit
            exit
            mesh-sdp 300:1 create
            exit
            stp
            exit
            no shutdown
      exit
      vpls 102 customer 1 m-vpls create
            spoke-sdp 102:2 create
               stp
                  path-cost 1
               exit
            exit
            mesh-sdp 300:2 create
            exit
            stp
            exit
            no shutdown
      exit
...
----------------------------------------------
*A:ALA-A2>config>service#
```

Configuring Selective MAC Flush

Use the following CLI syntax to enable selective MAC Flush in a VPLS.

CLI Syntax: `config>service# vpls` *`service-id`*
`send-flush-on-failure`

Use the following CLI syntax to disable selective MAC Flush in a VPLS.

CLI Syntax: `config>service# vpls` *`service-id`*
`no send-flush-on-failure`

ATM/Frame Relay PVC Access and Termination on a VPLS Service

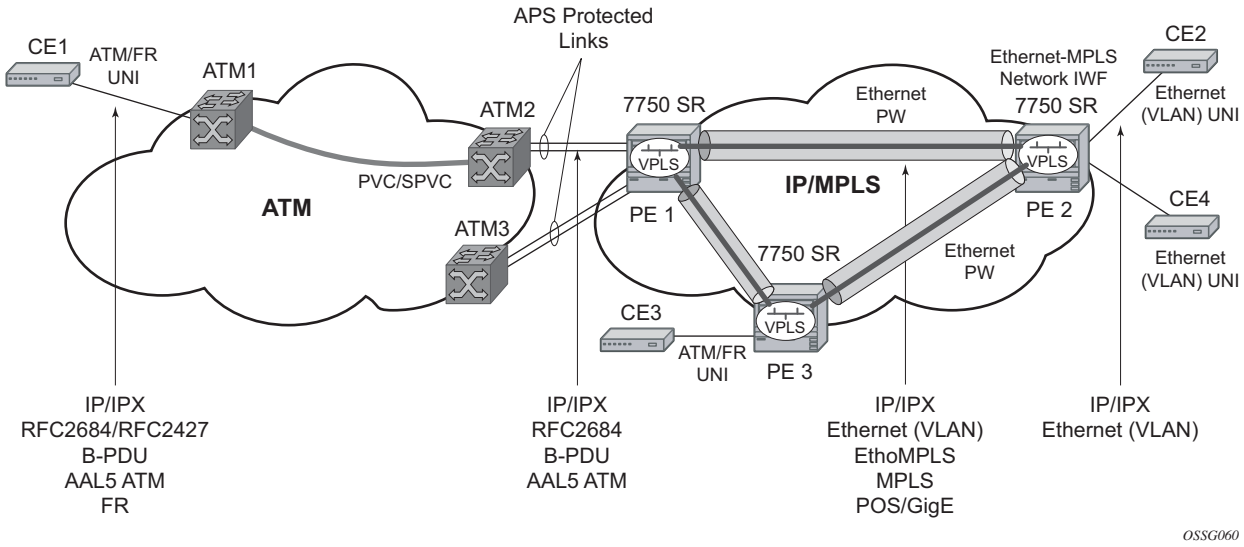

Figure 5: ATM/Frame Relay PVC Access and Termination on a VPLS Example

The application is depicted in Figure 47 provides access to a VPLS service to Frame Relay and ATM users connected either directly or through an ATM access network to a 7750 PE node. The 7750 SR supports a Frame Relay or an ATM VC-delimited Service Access Point (SAP) terminating on a VPLS service.

RFC 2427-encapsulated or RFC 2684-encapsulated untagged Ethernet/802.3 frames (with or without Frame Check Sequence (FCS)) or BPDUs from a customer's bridge device are received on a given SAP over an ATM or Frame Relay interface on the 7750 SR. The Frame Relay or ATM-related encapsulation is stripped and the frames (without FCS) are forwarded towards destination SAPs either locally, or using SDPs associated with the VPLS service (as dictated by destination MAC address VPLS processing). In the egress direction, the received untagged frames are encapsulated into RFC 2427 or RFC 2684 (no Q-tags are added, no FCS in the forwarded frame) and sent over ATM or a FR VC towards the customer CPE.

When AAL5 RFC2427/2684 encapsulated tagged frames are received from the customer's bridge on an FR/ATM SAP, the tags are transparent and the frames are processed as described above with the exception that the frames forwarded towards the destination(s) will have the received tags preserved. Similarly in the egress direction, the received tagged Ethernet frames are encapsulated as is (i.e. Q-tags are again transparent and preserved) into RFC 2427/2684 and sent over the FR/ATM PVC towards the customer CPE. Note that since the tagging is transparent, the 7750 SR performs unqualified MAC learning (i.e., MAC addresses are learned without reference to VLANs they are associated with). Because of that, MAC addresses used must be unique across all the VLANs used by the customer for a given VPLS service instance. If a customer wants a per-VLAN separation, then the VLAN traffic that needs to be separated must come on different VCs (different SAPs) associated with different VPLS service instances.

All VPLS functionality available on the 7750 SR is applicable to FR and ATM-delimited VPLS SAPs. For example, bridged PDUs received over ATM SAP can be tunneled through or dropped; all Forwarding Information Base functionality applies; packet level QoS and MAC filtering applies; etc. Also, split horizon groups are applicable to ATM SAPs terminating on VPLS. In other words, frame forwarding between ATM SAPs, also referred to as VCI-to-VCI forwarding, within the same group is disabled.

The Ethernet pseudowire is established using Targeted LDP (TLDP) signaling and uses the ether, vlan, or vpls VC type on the SDP. The SDP can be an MPLS or a GRE type.

Configuring Provider Edge Discovery Policies

Use the following CLI syntax to create PE discovery policy.

CLI Syntax:
```
config>service# pe-discovery-policy name
             password password
             polling-interval minutes
             server server-index address ip-address secret key
                [hash|hash2] [port port-num]
             timeout seconds
```

The following show an example of the PE discovery policy command usage.

Example:
```
config>service# pe-discovery-policy "RAD_Disc for Service 103"
config>service>pe-discovery-policy# password "timetravpn"
config>service>pe-discovery-policy# polling-interval 1
config>service>pe-discovery-policy# timeout 10
config>service>pe-discovery-policy# server 1 address 192.168.15.125
secret "LwyBQX4E2C/bXAGTtpNeYk" hash2 port 1812
config>service>pe-discovery-policy# server 2 address 192.168.15.122
secret "cj0n8F.5UUl5WBegZ.m6WmvwTYw6MZu0" hash2 port 1812
config>service>pe-discovery-policy# exit
```

The following displays the PE discovery policy configuration.

```
A:ALA-48>config>service# info
----------------------------------------------
        pe-discovery-policy "RAD_Disc for Service 103"
            password "timetravpn"
            polling-interval 1
            timeout 10
           server 1 address 192.168.15.125 secret "LwyBQX4E2C/bXAGTtpNeYk" hash2 port 1812
          server 2 address 192.168.15.122 secret "cj0n8F.5UUl5WBegZ.m6WmvwTYw6MZu0" hash2
port 1812
        exit
        customer 1 create
            description "Default customer"
        exit
...
----------------------------------------------
A:ALA-48
```

Configuring a VPLS Management Interface

A management interface behaves as a host (non-routing) similarly to how the out-of-band interfaces are created within the VPLS context. The following commands have the same definition as a regular interface. Notice that there may not be overlap in address space between ALL management interfaces regardless of service association.

This interface may be used for CPM protocols such as telnet, SSH, SNMP, ping, ANCP, etc. CPM filtering can be used to limit access to this interface.

Use the following CLI syntax to create a VPLS management interface.

CLI Syntax:
```
config>service>vpls# interface ip-int-name
address ip-address[/mask] [netmask]
arp-timeout seconds
description description-string
mac ieee-address
no shutdown
static-arp ip-address ieee-address
```

The following show an example of the PE discovery policy command usage.

Example:
```
config>service>vpls# interface test create
config>service>vpls>interface# address 123.231.10.10/24
config>service>vpls>interface# arp-timeout 5000
config>service>vpls>interface# mac 14:31:ff:00:00:00
config>service>vpls>interface# no shutdown
```

The following displays the configuration.

```
A:ALA-49>config>service>vpls>interface# info detail
----------------------------------------------
              no description
              mac 14:31:ff:00:00:00
              address 123.231.10.10/24
              no arp-timeout
              no shutdown
----------------------------------------------
A:ALA-49>config>service>vpls>interface#
```

Applying a PE Discovery Policy to a VPLS Service

Use the following CLI syntax to PE discovery parameters to a VPLS service.

CLI Syntax:
```
config>service# vpls service-id
            radius-discovery
               pe-discovery-policy name
               no shutdown
               user-name-format {vpn-id vpn-id | router-distinguisher rd}
            sap sap-id [split-horizon-group group-name]
               description description-string
            split-horizon-group group-name>[residential-group]
               restrict-protected-src [alarm-only]
```

Example:
```
config>service# vpls 103 customer 2 vpn 103 create
    config>service>vpls$ description "Default sap description for
service
    config>service>vpls$ split-horizon-group "SHG-RAD_Disc" create
    config>service>vpls>split-horizon-group$ restrict-protected-src
    config>service>vpls>split-horizon-group$ exit
    config>service>vpls# stp
    config>service>vpls>stp# no shutdown
    config>service>vpls>stp# exit
    config>service>vpls# radius-discovery
    config>service>vpls>radius-discovery# pe-discovery-policy "RAD_Disc
for Service 103"
    config>service>vpls>radius-discovery# user-name-format vpn-id
901:103
    config>service>vpls>radius-discovery# no shutdown
    config>service>vpls>radius-discovery# exit
    config>service>vpls# sap 1/1/7:0 create
    config>service>vpls>sap$ description "Default sap description for
service id 103"
    config>service>vpls>sap$ static-mac 12:34:56:78:90:0f create
    config>service>vpls>sap$ exit
    config>service>vpls# sap 1/1/11:0 create
    config>service>vpls>sap$ shutdown
    config>service>vpls>sap$ exit
    config>service>vpls# no sap 1/1/11:0
    config>service>vpls# sap 1/1/11:0 split-horizon-group "SHG-RAD_Disc"
create
    config>service>vpls>sap$ description "SHG for Radius Discovery"
    config>service>vpls>sap$ exit
    config>service>vpls# no shutdown
```

The following displays the VPLS PE discovery policy configuration.

```
*A:ALA-48>config>service>vpls# info
----------------------------------------------
            description "Default sap description for service id 103"
            split-horizon-group "SHG-RAD_Disc" create
                restrict-protected-src
            exit
            stp
                no shutdown
            exit
            radius-discovery
                pe-discovery-policy "RAD_Disc for Service 103"
                user-name-format vpn-id 901:103
                no shutdown
            exit
            sap 1/1/7:0 create
                description "Default sap description for service id 103"
                static-mac 12:34:56:78:90:0f create
            exit
            sap 1/1/11:0 split-horizon-group "SHG-RAD_Disc" create
                description "SHG for Radius Discovery"
            exit
            no shutdown
----------------------------------------------
*A:ALA-48>config>service>vpls#
```

Configuring Policy-Based Forwarding for Deep Packet Inspection in VPLS

The purpose policy-based forwarding is to capture traffic from a customer and perform a deep packet inspection (DPI) and forward traffic, if allowed, by the DPI.

In the following example, the split horizon groups are used to prevent flooding of traffic. Traffic from customers enter at SAP 1/1/5:5. Due to the mac-filter 100 that is applied on ingress, all traffic with dot1p 07 marking will be forwarded to SAP 1/1/22:1, which is the DPI.

DPI performs packet inspection/modifcation and either drops the traffic or forwards the traffic back into the box through SAP 1/1/21:1. Traffic will then be sent to spoke-sdp 3:5.

SAP 1/1/23:5 is configured to see if the VPLS service is flooding all the traffic. If flooding is peformed by the router then traffic would also be sent to SAP 1/1/23:5 (which it should not).

Figure 48 shows an example to configure policy-based forwarding for deep packet inspection on a VPLS service. For information about configuring filter policies, refer to the 7750 SR OS Router Configuration Guide.

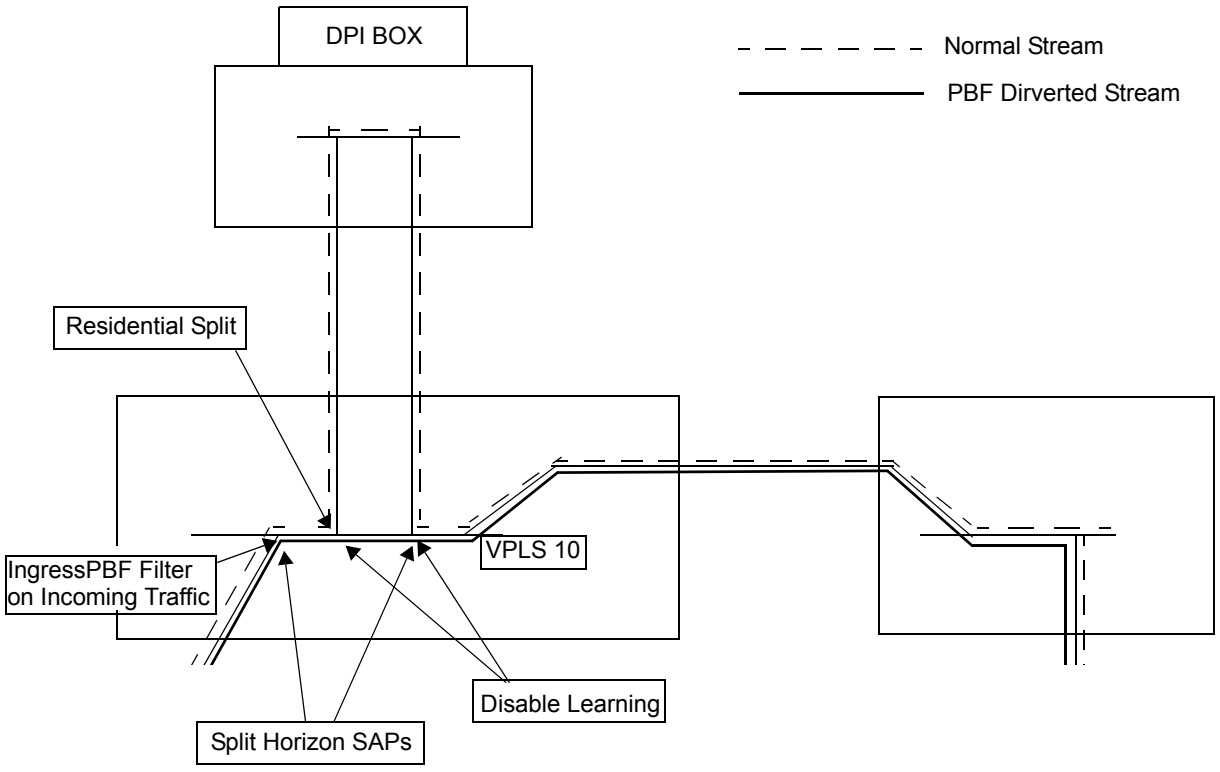

Figure 6: Policy-Based Forwarding For Deep Packet Inspection

Configuring the service:

Example:
```
config>service# vpls 10 customer 1 create
config>service>vpls$ service-mtu 1400
config>service>vpls$ split-horizon-group "dpi" residential-group create
config>service>vpls>split-horizon-group$ exit
config>service>vpls# split-horizon-group split create
config>service>vpls>split-horizon-group# exit
config>service>vpls# sap 1/1/21:1 split-horizon-group split create
config>service>vpls>sap$ disable-learning
config>service>vpls>sap$ static-mac 00:00:00:31:11:01 create
config>service>vpls>sap$ exit
config>service>vpls# sap 1/1/22:1 split-horizon-group "dpi" create
config>service>vpls>sap$ disable-learning
config>service>vpls>sap$ static-mac 00:00:00:31:12:01 create
config>service>vpls>sap$ exit
config>service>vpls# sap 1/1/23:5 create
config>service>vpls>sap$ static-mac 00:00:00:31:13:05 create
config>service>vpls>sap$ exit
config>service>vpls# no shutdown
```

The following example displays the service configuration:

```
*A:ALA-48>config>service# info
----------------------------------------------
...
        vpls 10 customer 1 create
            service-mtu 1400
            split-horizon-group "dpi" residential-group create
            exit
            split-horizon-group "split" create
            exit
            stp
                shutdown
            exit
            sap 1/1/21:1 split-horizon-group "split" create
                disable-learning
                static-mac 00:00:00:31:11:01 create
            exit
            sap 1/1/22:1 split-horizon-group "dpi" create
                disable-learning
                static-mac 00:00:00:31:12:01 create
            exit
            sap 1/1/23:5 create
                static-mac 00:00:00:31:13:05 create
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-48>config>service#
```

Configuring the MAC filter policy:

Example:
```
config>filter# mac-filter 100 create
config>filter>mac-filter$ default-action forward
config>filter>mac-filter$ entry 10 create
config>filter>mac-filter>entry$ match
config>filter>mac-filter>entry>match$ dot1p 07
config>filter>mac-filter>entry>match$ exit
config>filter>mac-filter>entry# log 101
config>filter>mac-filter>entry# action forward sap 1/1/22:1
config>filter>mac-filter>entry# exit
config>filter>mac-filter# exit
```

The following example displays the MAC filter configuration:

```
*A:ALA-48>config>filter# info
----------------------------------------------
...
        mac-filter 100 create
            default-action forward
            entry 10 create
                match
                    dot1p 7 7
                exit
                log 101
                action forward sap 1/1/22:1
            exit
        exit
...
----------------------------------------------
*A:ALA-48>config>filter#
```

Adding the MAC filter to the VPLS service:

Example:
```
config>service# config>service# vpls 10
config>service>vpls# sap 1/1/5:5 split-horizon-group "split" create
config>service>vpls>sap$ ingress
config>service>vpls>sap>ingress$ filter mac 100
config>service>vpls>sap>ingress$ exit
config>service>vpls>sap# static-mac 00:00:00:31:15:05 create
config>service>vpls>sap# exit
config>service>vpls# spoke-sdp 3:5 create
config>service>vpls>spoke-sdp$ exit
config>service>vpls# no shutdown
```

The following example displays the service configuration:

```
*A:ALA-48>config>service# info
----------------------------------------------
...
        vpls 10 customer 1 create
            service-mtu 1400
            split-horizon-group "dpi" residential-group create
            exit
            split-horizon-group "split" create
            exit
            stp
                shutdown
            exit
            sap 1/1/5:5 split-horizon-group "split" create
                ingress
                    filter mac 100
                exit
                static-mac 00:00:00:31:15:05 create
            exit
            sap 1/1/21:1 split-horizon-group "split" create
                disable-learning
                static-mac 00:00:00:31:11:01 create
            exit
            sap 1/1/22:1 split-horizon-group "dpi" create
                disable-learning
                static-mac 00:00:00:31:12:01 create
            exit
            sap 1/1/23:5 create
                static-mac 00:00:00:31:13:05 create
            exit
            spoke-sdp 3:5 create
            exit
            no shutdown
        exit
....
----------------------------------------------
*A:ALA-48>config>service#
```

Service Management Tasks

This section discusses the following service management tasks:

Modifying VPLS Service Parameters

You can change existing service parameters. The changes are applied immediately.
To display a list of services, use the `show service service-using vpls` command. Enter the parameter such as description, SAP, SDP, and/or service-MTU command syntax, and then enter the new information.

The following example displays command usage to modify VPLS parameters:

Example:
```
config>service# vpls 710
       config>service>vpls# description "This is a different
description"
       config>service>vpls# disable-aging
       config>service>vpls# disable-learning
       config>service>vpls# discard-unknown
       config>service>vpls# local-age 500
       config>service>vpls# remote-age 1000
```

The following displays the VPLS configuration modifications.

```
*A:ALA-1>config>service>vpls# info
----------------------------------------------
            description "This is a different description."
            disable-learning
            disable-aging
            discard-unknown
            local-age 500
            remote-age 1000
            stp
                shutdown
            exit
            sap 1/1/5:22 create
                description "VPLS SAP"
            exit
            spoke-sdp 2:22 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-1>config>service>vpls#
```

Modifying Management VPLS Parameters

To modify the range of VLANs on an access port that are to be managed by an existing management VPLS, first the new range should be entered and afterwards the old range removed. If the old range is removed before a new range is defined, all customer VPLS services in the old range will become unprotected and may be disabled.

CLI Syntax:
```
config>service# vpls service-id
     sap sap-id
        managed-vlan-list
           [no] range vlan-range
```

Example:
```
config>service# vpls 1
config>service>vpls# sap 1/1/1:1
config>service>vpls>sap# managed-vlan-list
config>service>vpls>vlan-list# range 100-800
config>service>vpls>vlan-list# no range 100-1000
config>service>vpls>vlan-list# exit
config>service>vpls>sap# exit
config>service>vpls# exit
```

Deleting a Management VPLS

As with normal VPLS service, a management VPLS cannot be deleted until SAPs and SDPs are unbound (deleted), interfaces are shutdown, and the service is shutdown on the service level.

Use the following CLI syntax to delete a management VPLS service:

CLI Syntax:
```
config>service
    [no] vpls service-id
        shutdown
        [no] mesh-sdp sdp-id
            shutdown
        [no] sap sap-id
            shutdown
```

Example:
```
config>service# vpls 1
config>service>vpls# sap 1/1/1:1
config>service>vpls>sap# shutdown
config>service>vpls>sap# exit
config>service>vpls# no sap 1/1/1:1
config>service>vpls# shutdown
config>service>vpls# exit
config>service# no vpls 1
```

Disabling a Management VPLS

You can shut down a management VPLS without deleting the service parameters.

When a management VPLS is disabled, all associated user VPLS services are also disabled (to prevent loops). If this is not desired, first un-manage the user's VPLS service by removing them from the managed-vlan-list or moving the spoke SDPs on to another tunnel SDP.

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

Example:
```
config>service# vpls 1
config>service>vpls# shutdown
config>service>vpls# exit
```

Deleting a VPLS Service

A VPLS service cannot be deleted until SAPs and SDPs are unbound (deleted), interfaces are shutdown, and the service is shutdown on the service level.

Use the following CLI syntax to delete a VPLS service:

CLI Syntax:
```
config>service
 [no] vpls service-id
    shutdown
    [no] mesh-sdp sdp-id
       shutdown
    sap sap-id [split-horizon-group group-name]
    no sap sap-id
       shutdown
```

Example:
```
config>service# vpls 10
config>service>vpls# mesh-sdp 6
config>service>vpls>mesh-sdp# shutdown
config>service>vpls>mesh-sdp# exit
config>service>vpls# no mesh-sdp 6
config>service>vpls# sap 1/1/5:22
config>service>vpls>sap# shutdown
config>service>vpls>sap# exit
config>service>vpls# no sap 1/1/5:22
config>service>vpls# shutdown
config>service>vpls# exit
config>service# no vpls 10
```

Disabling a VPLS Service

You can shut down an VPLS service without deleting the service parameters.

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

Example:
```
config>service# vpls 10
config>service>vpls# shutdown
config>service>vpls# exit
```

Re-enabling an VPLS Service

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

CLI Syntax: `config>service> vpls` *service-id*
`shutdown`

Example:
```
config>service# vpls 10
config>service>vpls# no shutdown
config>service>vpls# exit
```

VPLS Services Command Reference

Command Hierarchies

VPLS Service Configuration Commands

config
- **service**
  - [**no**] **vpls** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*] [**m-vpls**]
    - [**no**] **def-mesh-vc-id** *vc-id*
    - **description** *description-string*
    - **no description**
    - [**no**] **disable-aging**
    - [**no**] **disable-learning**
    - [**no**] **discard-unknown**
    - [**no**] **fdb-table-high-wmark** *high-water-mark*
    - [**no**] **fdb-table-low-wmark** *low-water-mark*
    - **fdb-table-size** *table-size*
    - **no fdb-table-size** [*table-size*]
    - **gsmp**
      - [**no**] **group** *name*
        - **ancp**
          - [**no**] **dynamic-topology-discover**
          - [**no**] **line-configuration**
          - [**no**] **oam**
        - **description** *description-string*
        - **no description**
        - **hold-multiplier** *multiplier*
        - **no hold-multiplier**
        - **keepalive** *seconds*
        - **no keepalive**
        - [**no**] **neighbor** *ip-address*
          - **description** *description-string*
          - **no description**
          - **local-address** *ip-address*
          - **no local-address**
          - **priority-marking dscp** *dscp-name*

```
                        — priority-marking prec ip-prec-value
                        — no priority-marking
                        — [no] shutdown
                    — [no] shutdown
                — [no] shutdown
— host-connectivity-verify source-ip ip-address [source-mac ieee-address] [interval
  interval] [action {remove|alarm}]
— igmp-snooping
        — mvr
                — description description-string
                — no description
                — group-policy policy-name
                — [no] shutdown
        — query-interval seconds
        — no query-interval
        — report-src-ip ip-address
        — no report-src-ip
        — robust-count robust-count
        — no robust-count
        — [no] shutdown
— [no] interface ip-int-name
        — [no] active-cpm-protocols
        — address ip-address[/mask]> [netmask]
        — no address
        — arp-timeout seconds
        — no arp-timeout
        — description description-string
        — no description
        — mac ieee-address
        — no mac
        — [no] shutdown
        — static-arp ip-address ieee-address
        — no static-arp ip-address [ieee-address]
— local-age aging-timer
— no local-age
— [no] mac-move
        — move-frequency frequency
        — no move-frequency
        — retry-timeout timeout
        — no retry-timeout
        — [no] shutdown
— mac-protect
        — [no] mac ieee-address
— mfib-table-high-wmark high-water-mark
— no mfib-table-high-wmark
— mfib-table-low-wmark low-water-mark
— no mfib-table-low-wmark
— mfib-table-size table-size
— no mfib-table-size
— [no] radius-discovery
        — pe-discovery-policy name
        — no pe-discovery-policy
        — [no] shutdown
        — user-name-format {vpn-id vpn-id | router-distinguisher rd}
        — no user-name-format
— remote-age aging-timer
— no remote-age
```

- [**no**] **send-flush-on-failure**
- **service-mtu** *octets*
- **no service-mtu**
- [**no**] **shutdown**
- [**no**] **split-horizon-group** *group-name* [*residential-group*]
  - **description** *description-string*
  - **no description**
  - **restrict-protected-src** *alarm-only*
  - **restrict-protected-src**
  - **restrict-unprotected-dst** *alarm-only*
  - **no restrict-unprotected-dst**
- **stp**
  - **forward-delay** *forward-delay*
  - **no forward-delay**
  - **hello-time** *hello-time*
  - **no hello-time**
  - **hold-count** *BDPU tx hold count*
  - **no hold-count**
  - **max-age** *max-info-age*
  - **no max-age**
  - **mode** {**rstp** | **comp-dot1w** | **dot1w** | **mst**}
  - **no mode**
  - [**no**] **mst-instance** *mst-inst-number*
    - **mst-priority** *bridge-priority*
    - **no mst-priority**
    - [**no**] **vlan-range** *vlan-range*
  - **mst-max-hops** *hops-count*
  - **no mst-max-hops**
  - **mst-name** *region-name*
  - **no mst-name**
  - **mst-revision** *revision-number*
  - **no mst-revision**
  - **priority** *bridge-priority*
  - **no priority**
  - [**no**] **shutdown**

SAP Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id] [vpn vpn-id] [mvpls]
            — sap sap-id [split-horizon-group group-name]
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy
                — anti-spoof {ip | mac | ip-mac}
                — no anti-spoof
                — arp-reply-agent [sub-ident]
                — no arp-reply-agent
                — atm
                    — egress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — encapsulation atm-encap-type
                    — ingress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — oam
                        — alarm-cells
                        — no alarm-cells
                — authentication-policy name
                — no authentication-policy
                — bpdu-translation {auto | pvst | stp}
                — no bpdu-translation
                — [no] collect-stats
                — description description-string
                — no description
                — dhcp
                    — description description-string
                    — no description
                    — lease-populate [nbr-of-entries]
                    — no lease-populate
                    — [no] option
                        — action [dhcp-action]
                        — no action
                        — circuit-id [ascii-tuple | vlan-ascii-tuple]
                        — [no] remote-id [mac | string string]
                        — [no] vendor-specific-option
                            — [no] client-mac-address
                            — [no] sap-id
                            — [no] service-id
                            — string text
                            — no string
                            — [no] system-id
                    — proxy-server
                        — emulated-server ip-address
                        — no emulated-server
                        — lease-time [days days] [hrs hours] [min minutes] [sec
                          seconds] [radius-override]
                        — no lease-time
                        — [no] shutdown
```

```
            — [no] shutdown
            — [no] snoop
      — [no] disable-aging
      — [no] disable-learning
      — [no] discard-unknown-source
      — egress
            — agg-rate-limit agg-rate
            — no agg-rate-limit agg-r
            — filter ip ip-filter-id
            — filter ipv6 ipv6-filter-id
            — filter mac mac-filter-id
            — no filter
            — no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
            — multicast-group group-name
            — no multicast-group
            — [no] qinq-mark-top-only
            — qos policy-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
                        — no mbs
                        — rate pir-rate [cir cir-rate]
                        — no rate
            — [no] scheduler-override
                  — [no] scheduler scheduler-name
                        — rate pir-rate [cir cir-rate]
                        — no rate
            — scheduler-policy scheduler-policy-name
            — no scheduler-policy
      — host {[ip ip-address] [mac ieee-address]}[subscriber sub-ident-string]
        [sub-profile sub-profile-name] [sla-profile sla-profile-name] [ancp-string
        ancp-string]
      — no host {[ip ip-address] [mac ieee-address]}
      — no host all
      — host-connectivity-verify source-ip ip-address [source-mac ieee-address]
        [interval interval] [action {remove | alarm}]
      — igmp-snooping
            — [no] fast-leave
            — import policy-name
            — no import
            — last-member-query-interval interval
            — no last-member-query-interval
            — max-num-groups max-num-groups
            — no max-num-groups
            — mcac
```

```
                    — mc-constraints
                        — level level-id bw bandwidth
                        — no  level level-id
                        — number-down number-lag-port-down level level-id
                        — no number-down
                    — policy policy-name
                    — no policy
                    — unconstrained-bw bandwidth mandatory-bw mandatory-bw
                    — no unconstrained-bw
                — [no] mrouter-port
                — mvr
                    — from-vpls vpls-id
                    — no from-vpls
                    — to-sap sap-id
                    — no to-sap
                — query-interval interval
                — no query-interval
                — query-response-interval interval
                — no query-response-interval
                — robust-count count
                — no robust-count
                — [no] send-queries
                — static
                    — [no] group group-address
                        — [no] source ip-address
                        — [no] starg
                — version version
                — no version
            — ingress
                — filter ip ip-filter-id
                — filter ipv6 ipv6-filter-id
                — filter mac mac-filter-id
                — no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
                — match-qinq-dot1p {top | bottom}
                — no match-qinq-dot1p
                — qos policy-id [shared-queuing | multipoint-shared]
                — no qos
                — [no] queue-override
                    — [no] queue queue-id
                        — adaptation-rule  [pir {max|min|closest}] [cir {max | min | closest}]
                        — no adaptation-rule
                        — cbs size-in-kbytes
                        — no cbs
                        — high-prio-only percent
                        — no high-prio-only
                        — mbs  size-in-kbytes
                        — no mbs
                        — rate pir-rate [cir cir-rate]
                        — no rate
                — [no] scheduler-override
                    — [no] scheduler scheduler-name
                        — rate pir-rate [cir cir-rate]
                        — no rate
                — scheduler-policy scheduler-policy-name
```

- — **no scheduler-policy**
- — [**no**] **l2pt-termination**
- — **limit-mac-move** [**blockable** | **non-blockable**]
- — **no limit-mac-move**
- — [**no**] **mac-pinning**
- — **managed-vlan-list**
  - — [**no**] **default-sap**
  - — [**no**] **range** *vlan-range*
- — **max-nbr-mac-addr** *table-size*
- — **no max-nbr-mac-addr**
- — **multi-service-site** *customer-site-name*
- — **no multi-service-site**
- — **restrict-protected-src** *alarm-only*
- — **no restrict-protected-src**
- — **restrict-unprotected-dst** *alarm-only*
- — **no restrict-unprotected-dst**
- — [**no**] **shutdown**
- — [**no**] **static-mac** *ieee-address*
- — **stp**
  - — [**no**] **auto-edge**
  - — [**no**] **edge-port**
  - — **link-type** {**pt-pt** | **shared**}
  - — **no link-type** [**pt-pt** | **shared**]
  - — **mst-instance** *mst-inst-number*
    - — **mst-path-cost** *inst-path-cost*
    - — **no mst-path-cost**
    - — **mst-port-priority** *stp-priority*
    - — **no mst-port-priority**
  - — **path-cost** *sap-path-cost*
  - — **no path-cost**
  - — [**no**] **port-num** *virtual-port-number*
  - — **port-num** *stp-priority*
  - — **no port-num**
  - — **priority** *stp-priority*
  - — **no priority**
  - — [**no**] **root-guard**
  - — [**no**] **shutdown**
- — [**no**] **sub-sla-mgmt**
  - — **def-sla-profile** *default-sla-profile-name*
  - — **no def-sla-profile**
  - — **def-sub-profile** *default-subscriber-profile-name*
  - — **no def-sub-profile**
  - — [**no**] **mac-da-hashing**
  - — **multi-sub-sap** [*subscriber-limit*]
  - — [**no**] **shutdown**
  - — **single-sub-parameters**
    - — **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*]
    - — **no non-sub-traffic**
    - — [**no**] **profiled-traffic-only**
  - — **sub-ident-policy** *sub-ident-policy-name*
  - — **no sub-ident-policy**
- — **tod-suite** *tod-suite-name*
- — **no tod-suite**

Mesh SDP Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id] [vpn vpn-id] [mvpls]
            — mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
            — no mesh-sdp sdp-id[:vc-id]
                — accounting-policy acct-policy-id
                — no accounting-policy
                — [no] collect-stats
                — dhcp
                    — description description-string
                    — no description
                    — snoop [l2-header]
                    — no snoop
                — egress
                    — filter ip ip-filter-id
                    — filter ipv6 ipv6-filter-id
                    — filter mac mac-filter-id
                    — no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
                    — mfib-allowed-mda-destinations
                        — [no] mda mda-id
                    — vc-label egress-vc-label
                    — no vc-label [egress-vc-label]
                — igmp-snooping
                    — [no] fast-leave
                    — import policy-name
                    — no import
                    — last-member-query-interval interval
                    — no last-member-query-interval
                    — max-num-groups max-num-groups
                    — no max-num-groups
                    — mcac
                        — policy policy-name
                        — no policy
                        — unconstrained-bw bandwidth mandatory-bw mandatory-bw
                        — no unconstrained-bw
                    — [no] mrouter-port
                    — query-interval interval
                    — no query-interval
                    — query-response-interval interval
                    — no query-response-interval
                    — robust-count count
                    — no robust-count
                    — [no] send-queries
                    — static
                        — [no] group group-address
                            — [no] source ip-address
                            — [no] starg
                    — version version
                    — no version
                — ingress
                    — filter ip ip-filter-id
                    — filter ipv6 ipv6-filter-id
                    — filter mac mac-filter-id
                    — no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
```

— **vc-label** *ingress-vc-label*
— **no vc-label** [*ingress-vc-label*]
— [**no**] **mac-pinning**
— [**no**] **shutdown**
— [**no**] **static-mac** ***ieee-address***
— **vlan-vc-tag** *0..4094*
— **no vlan-vc-tag** [*0..4094*]

Spoke SDP Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id] [vpn vpn-id] [mvpls]
            — spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [split-horizon-group group-name]
            — no spoke-sdp sdp-id[:vc-id]
                — accounting-policy acct-policy-id
                — no accounting-policy
                — bpdu-translation {auto | pvst | stp}
                — no bpdu-translation
                — [no] collect-stats
                — dhcp
                    — description description-string
                    — no description
                    — [no] snoop
                — [no] disable-aging
                — [no] disable-learning
                — [no] discard-unknown-source
                — egress
                    — filter ip ip-filter-id
                    — filter ipv6 ipv6-filter-id
                    — filter mac mac-filter-id
                    — no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
                    — mfib-allowed-mda-destinations
                        — [no] mda mda-id
                    — vc-label egress-vc-label
                    — no vc-label [egress-vc-label]
                — igmp-snooping
                    — [no] fast-leave
                    — import policy-name
                    — no import
                    — last-member-query-interval interval
                    — no last-member-query-interval
                    — max-num-groups max-num-groups
                    — no max-num-groups
                    — mcac
                        — policy policy-name
                        — no policy
                        — unconstrained-bw bandwidth mandatory-bw mandatory-bw
                        — no unconstrained-bw
                    — [no] mrouter-port
                    — query-interval interval
                    — no query-interval
                    — query-response-interval interval
                    — no query-response-interval
                    — robust-count count
                    — no robust-count
                    — [no] send-queries
                    — static
                        — [no] group group-address
                            — [no] source ip-address
                            — [no] starg
```

- **version** *version*
- **no version**

— **ingress**

- **filter ip** *ip-filter-id*
- **filter ipv6** *ipv6-filter-id*
- **filter mac** *mac-filter-id*
- **no filter** [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]
- **vc-label** *ingress-vc-label*
- **no vc-label** [*ingress-vc-label*]

— [**no**] **l2pt-termination**
— **limit-mac-move** [**blockable** | **non-blockable**]
— **no limit-mac-move**
— [**no**] **mac-pinning**
— **max-nbr-mac-addr** *table-size*
— **no max-nbr-mac-addr**
— [**no**] **shutdown**
— [**no**] **static-mac** *ieee-address*
— **stp**

- [**no**] **auto-edge**
- [**no**] **edge-port**
- **link-type** {**pt-pt** | **shared**}
- **no link-type** [**pt-pt** | **shared**]
- **path-cost** *sap-path-cost*
- **no path-cost**
- [**no**] **port-num** *virtual-port-number*
- **priority** *stp-priority*
- **no priority**
- [**no**] **root-guard**
- [**no**] **shutdown**

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

Egress Multicast Group Commands

```
config
    — service
        — egress-multicast-group group-name [create]
        — no egress-multicast-group group-name
            — description description-string
            — no description
            — dest-chain-limit destinations per pass
            — no dest-chain-limit
            — sap-common-requirements
                — dot1q-etype 0x0600..0xffff
                — no dot1q-etype
                — egress-filter [ip ip-filter-id]
                — egress-filter [ipv6 ipv6-filter-id]
                — egress-filter [mac mac-filter-id]
                — no egress-filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
                — encap-type {dot1q | null}
                — no encap-type
config
    — service
        — [no] vpls service-id [customer customer-id] [vpn vpn-id] [mvpls]
            — sap sap-id [split-horizon-group group-name]
            — no sap sap-id
                — egress
                    — multicast-group group-name
                    — no multicast-group
```

Provider Edge Discovery Policy Commands

```
config
    — service
        — [no] pe-discovery-policy name
            — password password
            — no password
            — polling-interval minutes
            — no polling-interval
            — server server-index address ip-address secret key [hash | hash2] [port port-num]
            — no server server-index
            — timeout seconds
            — no timeout
```

Show Commands

show
- **service**
  - **egress-label** *egress-label1* [*egress-label2*]
  - **fdb-info**
  - **fdb-mac** *ieee-address* [**expiry**]
  - **id** *service-id*
    - **all**
    - **authentication**
      - **statistics** [**policy** *name*] [**sap** *sap-id*]
    - **base**
    - **dhcp**
      - **lease-state** [[**sap** *sap-id*] | [**sdp** *sdp-id*:*vc-id*] | [**interface** *interface-name*] | [**ip-address** *ip-address*]] [**detail**] |[**mac** *ieee-address*] |[**wholesaler** *service-id*][**detail**]
      - **statistics** [**sap** *sap-id*
      - **statistics** [**sdp** *sdp-id*:*vc-id*]
      - **statistics** [**interface** *interface-name*]
      - **summary**
    - **fdb** [**sap** *sap-id*] [**expiry**]] | [**sdp** *sdp-id* [**expiry**]] | [**mac** *ieee-address* [**expiry**]] | [**detail**][**expiry**]
    - **gsmp**
      - **neighbors group** [*name*] [*ip-address*]
      - **sessions** [**group** *name*] **neighbor** *ip-address*] [**port** *port-number*] [**association**] [**statistics**]
    - **host** [**sap** *sap-id*] [**detail**]
    - **host summary**
    - **egress-multicast-group statistics** [**sap** *sap-id*]
    - **igmp-snooping**
      - **all**
      - **base**
      - **mrouters** [**detail**]
      - **mvr**
      - **port-db** {**sap** *sap-id* | **sdp** *sdp-id*:*vc-id*} [**group** *grp-address*] | **detail**]
      - **proxy-db** [**group** *grp-address* | **detail**]
      - **querier**
      - **static** [**sap** *sap-id* | **sdp** *sdp-id*:*vc-id*]
      - **statistics** [**sap** *sap-id* | **sdp** *sdp-id*:*vc-id*]
    - **labels**
    - **mac-protect**
    - **mfib** [**brief** | **group** *grp-address* | **statistics** [**group** *grp-address*]]
    - **mstp-configuration**
    - **retailers**
    - **sap** [*sap-id* [**detail**]]
    - **sdp** [*sdp-id* | **far-end** *ip-addr*] [**detail**]
    - **split-horizon-group** [*group-name*]
    - **stp** [**detail**]
    - **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-address*[/*mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**detail**]
    - **wholesalers**
  - **ingress-label** *start-label* [*end-label*]
  - **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
                    — sdp [sdp-id | far-end ip-address] [detail | keep-alive-history]
                    — sdp-using [sdp-id[:vc-id] | far-end ip-address]
                    — sdp [sdp-id | far-end ip-addr] [detail | keep-alive-history]
                    — sdp-using [sdp-id[:vc-id] | far-end ip-address]
                    — service-using [vpls] [sdp sdp-id] [customer customer-id]
                    — subscriber-using [service-id service-id] [sap-id sap-id] [interface ip-int-name] [ip ip-
                      address[/mask]] [mac ieee-address] [sub-profile sub-profile-name] [sla-profile sla-profile-
                      name]
show
    — egress-replication
                    — vpls vpls-service-id  mda slot/mda
                    — vpls vpls-service-id  mda slot/mda [igmp-record grp-address {source source-ip-address |
                      starg}] | [mRouter]
show
    — igmp
                    — group [grp-ip-address]
                    — ssm-translate
                    — interface  [ip-int-name | ip-address] [group grp-address] [detail]
                    — static [ip-int-name | ip-addr]
                    — statistics [ip-int-name | ip-address]
                    — status
```

Clear Commands

```
clear
    — service
                    — id service-id
                                — authentication
                                            — statistics
                                — dhcp
                                            — lease-state  [no-dhcp-release]
                                            — lease-state ip-address ip-address[/mask] [no-dhcp-release]
                                            — lease-state mac ieee-address[no-dhcp-release]
                                            — lease-state sap sap-id [no-dhcp-release]
                                            — lease-state sdp sdp-id:vc-id [no-dhcp-release]
                                            — statistics [sap sap-id | sdp  [sdp-id[:vc-id] ]
                                — fdb {all | mac ieee-address | sap sap-id | mesh-sdp sdp-id[:vc-id] | spoke-sdp sdp-
                                  id:vc-id}
                                — igmp-snooping
                                            — port-db sap sap-id [group grp-address [source ip-address]]
                                            — port-db sdp sdp-id:vc-id [group grp-address [source ip-address]]
                                            — querier
                                            — statistics [all | sap sap-id | sdp sdp-id:vc-id
                                — mfib
                                            — statistics [all | group grp-address ]
                                — mesh-sdp sdp-id[:vc-id] ingress-vc-label
                                — spoke-sdp sdp-id:vc-id ingress-vc-label
                                — stp
                                            — detected-protocols [ all |  sap sap-id |  spoke-sdp [sdp-id[:vc-id] ]
                    — statistics
                                — id service-id
```

```
                        — counters
                        — sap sap-id {all | counters | stp}
                        — spoke-sdp sdp-id[:vc-id] {all | counters | stp}
                        — stp
                    — sap sap-id {all | counters | stp}
                    — sdp sap-id {keep-alive}

clear
    — router
        — dhcp
            — statistics [interface ip-int-name | ip-address]
```

Debug Commands

```
debug
    — service
        — id service-id
        — stp
            — all-events
            — [no] bpdu
            — [no] core-connectivity
            — [no] exception
            — [no] fsm-state-changes
            — [no] fsm-timers
            — [no] port-role
            — [no] port-state
            — [no] sap sap-id
            — [no] sdp sdp-id:vc-id
debug
    — igmp
        — router
            — [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]
```

VPLS Service Configuration Commands

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>service>vpls
config>service>vpls>gsmp
config>service>vpls>gsmp>group
config>service>vpls>gsmp>group>neighbor
config>service>vpls>interface
config>service>vpls>split-horizon-group
config>service>vpls>sap
config>service>vpls>sap>stp
config>service>vpls>sap>sub-sla-mgmt
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>vpls>spoke-sdp>stp
config>service>vpls>stp
config>service>vpls>sap>dhcp>proxy
config>service>vpls>radius-discovery

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 **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.

Service Operational State — A service is regarded as operational providing that two SAPs or if one SDP are operational.

SDP (global) — When an SDP is shutdown at the global service level, all bindings to that SDP are put into the out-of-service state and the SDP itself is put into the administratively and operationally down states. Packets that would normally be transmitted using this SDP binding will be discarded and counted as dropped packets.

SDP (service level) — Shutting down an SDP within a service only affects traffic on that service from entering or being received from the SDP. The SDP itself may still be operationally up for other services.

SDP Keepalives — Enables SDP connectivity monitoring keepalive messages for the SDP ID. Default state is disabled (shutdown) in which case the operational state of the SDP-ID is not affected by the keepalive message state.

VPLS SAPs and SDPs — SAPs are created in a VPLS and SDPs are bound to a VPLS in the administratively up default state. The created SAP will attempt to enter the operationally up state. An SDP will attempt to go into the in-service state once bound to the VPLS.

description

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

Context config>service>vpls
config>service>vpls>gsmp>group
config>service>vpls>gsmp>group>neighbor
config>service>vpls>igmp-snooping>mvr
config>service>vpls>interface
config>service>vpls>split-horizon-group
config>service>vpls>sap
config>service>vpls>sap>dhcp

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

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

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

Default No description associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters 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.

VPLS Service Commands

vpls

Syntax **vpls** *service-id* **customer** *customer-id* **vpn** *vpn-id* [**m-vpls**]
vpls *service-id*
no vpls *service-id*

Context config>service

Description This command creates or edits a Virtual Private LAN Services (VPLS) instance.

The **vpls** command is used to create or maintain a VPLS service. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

A VPLS service connects multiple customer sites together acting like a zero-hop, layer 2 switched domain. A VPLS is always a logical full mesh.

When a service is created, the **create** keyword must be specified if the **create** command is enabled in the **environment** context.

When a service is created, the **customer** keyword and *customer-id* must be specified and 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.

More than one VPLS service may be created for a single customer ID.

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

The **no** form of this command deletes the VPLS service instance with the specified *service-id*. The service cannot be deleted until all SAPs and SDPs defined within the service ID have been shutdown and deleted, and the service has been shutdown.

service-id — The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR on which this service is defined.

Values 1 — 2147483647

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)

m-vpls — Specifies a managed VPLS.

bpdu-translation

Syntax **bpdu-translation** {**auto** | **pvst** | **stp**}
no bpdu-translation

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command enables the translation of BPDUs to a given format, meaning that all BPDUs transmitted on a given SAP or spoke SDP will have a specified format.

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

Default **no bpdu-translation**

Parameters **auto —** Specifies that appropriate format will be detected automatically, based on type of bpdus received on such port.

pvst — Specifies the BPDU-format as PVST. Note that the correct VLAN tag is included in the payload (depending on encapsulation value of outgoing SAP).

stp — Specifies the BPDU-format as STP.

l2pt-termination

Syntax [**no**] **l2pt-termination**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This commands enables Layer 2 Protocol Tunneling (L2PT) termination on a given SAP or spoke SDP. L2TP termination will be supported only for STP BPDUs. PDUs of other protocols will be discarded.

This feature can be enabled only if STP is disabled in the context of the given VPLS service.

Default **no l2pt-termination**

def-mesh-vc-id

Syntax [**no**] **def-mesh-vc-id** *vc-id*

Context config>service>vpls

Description This command configures the value used by each end of a tunnel to identify the VC. If this command is not configured, then the service ID value is used as the VC-ID.

This VC-ID is used instead of a label to identify a virtual circuit.The VC-ID is significant between peer SRs on the same hierarchical level. The value of a VC-ID is conceptually independent from the value of the label or any other datalink specific information of the VC.

The **no** form of this command disables the VC-ID.

Default **none**

Values 1 — 4294967295

disable-aging

Syntax [**no**] **disable-aging**

Context config>service>vpls
config>service>vpls>spoke-sdp
config>service>vpls>sap

Description This command disables MAC address aging across a VPLS service or on a VPLS service SAP or spoke SDP.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the VPLS forwarding database (FDB). The **disable-aging** command turns off aging for local and remote learned MAC addresses.

When **no disable-aging** is specified for a VPLS, it is possible to disable aging for specifc SAPs and/or spoke SDPs by entering the **disable-aging** command at the appropriate level.

When the **disable-aging** command is entered at the VPLS level, the **disable-aging** state of individual SAPs or SDPs will be ignored.

The **no** form of this command enables aging on the VPLS service.

Default **no disable-aging**

disable-learning

Syntax **disable-learning**
no disable-learning

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command enables learning of new MAC addresses in the VPLS forwarding database (FDB) for the service instance, SAP instance or spoke SDP instance.

When **disable-learning** is enabled, new source MAC addresses will not be entered in the VPLS service forwarding database. This is true for both local and remote MAC addresses.

When **disable-learning** is disabled, new source MAC addresses will be learned and entered into the VPLS forwarding database.

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

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

Default **no disable-learning (Normal MAC learning is enabled)**

discard-unknown

Syntax [**no**] **discard-unknown**

Context config>service>vpls

Description By default, packets with unknown destination MAC addresses are flooded. If discard-unknown is enabled at the VPLS level, packets with unknown destination MAC address will be dropped instead (even when configured FIB size limits for VPLS or SAP are not yet reached).

The **no** form of this command allows flooding of packets with unknown destination MAC addresses in the VPLS.

Default **no discard-unknown** - Packets with unknown destination MAC addresses are flooded

fdb-table-high-wmark

Syntax [**no**] **fdb-table-high-wmark** *high-water-mark*

Context config>service>vpls

Description This command specifies the value to send logs and traps when the threshold is reached.

Parameters *high-water-mark* — Specify the value to send logs and traps when the threshold is reached

Values 1 — 100

Default 95%

fdb-table-low-wmark

Syntax [**no**] **fdb-table-low-wmark** *low-water-mark*

Context config>service>vpls

Description This command specifies the value to send logs and traps when the threshold is reached.

Parameters *low-water-mark* — Specify the value to send logs and traps when the threshold is reached.

Values 1 — 100

Default 90%

fdb-table-size

Syntax **fdb-table-size** *table-size*
no fdb-table-size [*table-size*]

Context config>service>vpls

Description This command specifies the maximum number of MAC entries in the forwarding database (FDB) for the VPLS instance on this node.

The **fdb-table-size** specifies the maximum number of forwarding database entries for both learned and static MAC addresses for the VPLS instance.

The **no** form of this command returns the maxium FDB table size to default.

Default **250** — Forwarding table of 250 MAC entries

Values 1 — 196607
Chassis-mode A or B limit: 131071
Chassis-mode C limit: 196607

interface

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

Context config>service>vpls

Description This command creates an IP interface.

active-cpm-protocols

Syntax [**no**] **active-cpm-protocols**

Context config>service>vpls>interface

Description This command enables CPM protocols on this interface.

address

Syntax **address** *ip-address*[*/mask*]> [*netmask*]
no address

Context config>service>vpls>interface

Description This command assigns an IP address, IP subnet, and broadcast address format to an IES IP router interface. Only one IP address can be associated with an IP interface.

An IP address must be assigned to each IES IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the 7750 SR.

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

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

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

Use the **no** form of this command to remove the IP address assignment from the IP interface. When the **no address** command is entered, the interface becomes operationally down.

Address	Admin state	Oper state
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an adminstratively 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 netmask

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.

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>vpls>interface

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.

When the **arp-populate** and **lease-populate** commands are enabled on an IES interface, the ARP table entries will no longer be dynamically learned, but instead by snooping DHCP ACK message from a DHCP server. In this case the configured **arp-timeout** value has no effect.

The default value for **arp-timeout** is 14400 seconds (4 hours).

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

mac

Syntax **mac** *ieee-address*
no mac

Context config>service>vpls>interface

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

For Routed Central Office (CO), a group interface has no IP address explicitely 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.

static-arp

Syntax **static-arp** *ip-address ieee-address*
no static-arp *ip-address* [*ieee-address*]

Context config>service>vpls>interface

Description This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP 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.

local-age

Syntax **local-age** *seconds*
no local-age

Context config>service>vpls

Description Specifies the aging time for locally learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance.

In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the FDB. The **local-age** timer specifies the aging time for local learned MAC addresses.

The **no** form of this command returns the local aging timer to the default value.

Default **local age 300** — local MACs aged after 300 seconds.

Parameters *seconds* — The aging time for local MACs expressed in seconds.

Values 60 — 86400

mac-move

Syntax [**no**] **mac-move**

Context config>service>vpls

Description This command enables the context to configure MAC move attributes. A sustained high re-learn rate can be a sign of a loop somewhere in the VPLS topology. Typically, STP detects loops in the topology, but for those networks that do not run STP, the mac-move feature is an alternative way to protect your network against loops.

When enabled in a VPLS, **mac-move** monitors the re-learn rate of each MAC. If the rate exceeds the configured maximum allowed limit, it disables the SAP where the source MAC was last seen. The SAP can be disabled permanently (until a **shutdown**/**no shutdown** command is executed) or for a length of time that grows linearly with the number of times the given SAP was disabled. You have the option of marking a SAP as non-blockable in the **config>service>vpls>sap>limit-mac-move** or **config>service>vpls>spoke-sdp>limit-mac-move** contexts, see page 497. This means that when the re-learn rate has exceeded the limit, another (blockable) SAP will be disabled instead.

The **mac-move** command enables the feature at the service level for SAPs and spoke SDPs, as only those objects can be blocked by this feature. Mesh SDPs are never blocked, but their re-learn rates (sap-to-mesh/spoke-to-mesh or vice versa) are still measured.

The operation of this feature is the same on the SAP and spoke SDP. For example, if a MAC address moves from SAP to SAP, from SAP to spoke SDP, or between spoke SDPs, one will be blocked to prevent thrashing. If the MAC address moves between a SAP and mesh SDP or spoke SDP and mesh SDP combinations, the respective SAP or spoke SDP will be blocked.

The re-learn rate is computed as the number of times a MAC moves in a 5 second interval. Therefore, the fastest a loop can be detected and broken is 5 seconds.

The **no** form of this command disables MAC move.

Default **not enabled**

mac-protect

Syntax **mac-protect**

Context config>service>vpls

Description This command indicates whether or not this MAC is protected. When enabled, the agent will protect the MAC from being learned or re-learned on a SAP that has restricted learning enabled.

Default **disabled**

Sample Output

```
A:ALA-48# show service id <service-id> mac-protect
===============================================================================
Mac Protection
===============================================================================
ServId    MAC
-------------------------------------------------------------------------------
1         aa:aa:aa:aa:aa:ab
-------------------------------------------------------------------------------
No. of MAC Entries: 1
===============================================================================
A:ALA-48#
```

mac

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

Context config>service>vpls>mac-protect

Description This command adds a protected MAC address entry.

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

move-frequency

Syntax **move-frequency** *frequency*
no move-frequency

Context config>service>vpls>mac-move

Description This indicates the maximum rate at which MAC's can be re-learned in the VPLS service, before the SAP where the moving MAC was last seen is automatically disabled in order to protect the system against undetected loops or duplicate MAC's.

The rate is computed as the maximum number of re-learns allowed in a 5 second interval. For example, the default rate of 10 relearns per second corresponds to 50 relearns in a 5 second period.

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

Default **2** (when **mac-move** is enabled)

Parameters *frequency* — Specifies the rate, in 5-second intervals for the maximum number of relearns.

Values 1 — 100

retry-timeout

Syntax **retry-timeout** *timeout*
no retry-timeout

Context config>service>vpls>mac-move

Description This indicates the time in seconds to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled.

A zero value indicates that the SAP will not be automatically re-enabled after being disabled. If, after the SAP is reenabled it is disabled again, the effective retry timeout is doubled in order to avoid thrashing.

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

Default **10** (when **mac-move** is enabled)

Parameters *timeout* — Specifies the time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled.

Values 0 — 120

mfib-table-high-wmark

Syntax [**no**] **mfib-table-high-wmark** *high-water-mark*

Context config>service>vpls

Description This command specifies the multicast FIB high watermark. When the percentage filling level of the multicast FIB exceeds the configured value, a trap is generated and/or a log entry is added.

Parameters *high-water-mark* — Specifies the multicast FIB high watermark as a percentage.

Values 1 — 100

Default 95%

mfib-table-low-wmark

Syntax [**no**] **mfib-table-low-wmark** *low-water-mark*

Context config>service>vpls

Description This command specifies the multicast FIB low watermark. When the percentage filling level of the Multicast FIB drops below the configured value, the corresponding trap is cleared and/or a log entry is added.

Parameters *low-water-mark* — Specifies the multicast FIB low watermark as a percentage.

Values 1 — 100

Default 90%

mfib-table-size

Syntax **mfib-table-size** *size*
no mfib-table-size

Context config>service>vpls

Description This command specifies the maximum number of (s,g) entries in the multicast forwarding database (MFIB) for this VPLS instance.

The *mfib-table-size* parameter specifies the maximum number of multicast database entries for both learned and static multicast addresses for the VPLS instance.

When a table-size limit is set on the mfib of a service which is lower than the current number of dynamic entries present in the mfib then the number of entries remains above the limit.

The **no** form of this command removes the configured maxium MFIB table size.

Default **none**

Parameters *size* — The maximum number of (s,g) entries allowed in the Multicast FIB.

Values 1 — 16383

remote-age

Syntax **remote-age** *seconds*
no remote-age

Context config>service>vpls

Description Specifies the aging time for remotely learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance.

In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Like in a layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the FDB. The **remote-age** timer specifies the aging time for remote learned MAC addresses. To reduce the amount of signaling required between switches configure this timer larger than the **local-age** timer.

The **no** form of this command returns the remote aging timer to the default value.

Default **remote age 900** — Remote MACs aged after 900 seconds

Parameters *seconds* — The aging time for remote MACs expressed in seconds.

Values 60 — 86400

send-flush-on-failure

Syntax [**no**] **send-flush-on-failure**

Context config>service>vpls

Description This command enables sending out “flush-all-from-ME” messages to all LDP peers included in affected VPLS, in the event of physical port failures or “oper-down” events of individual SAPs. This feature provides an LDP-based mechanism for recovering a physical link failure in a dual-homed connection to a VPLS service. This method provides an alternative to RSTP solutions where dual homing redundancy and recovery, in the case of link failure, is resolved by RSTP running between a PE router and CE devices.

This feature cannot be enabled on management VPLS.

Default **no send-flush-on-failure**

service-mtu

Syntax **service-mtu** *octets*
no service-mtu

Context config>service>vpls

Description This command configures the service payload (Maximum Transmission Unit – MTU), in bytes, for the service. This MTU value overrides the service-type default MTU.

The **service-mtu** defines the payload capabilities of the service. It is used by the system to validate the SAP and SDP binding’s operational state within the service.

The service MTU and a SAP’s service delineation encapsulation overhead (i.e., 4 bytes for a dot1q tag) is used to derive the required MTU of the physical port or channel on which the SAP was created. If the required payload is larger than the port or channel MTU, then the SAP will be placed in an inoperative state. If the required MTU is equal to or less than the port or channel MTU, the SAP will be able to transition to the operative state.

When binding an SDP to a service, the service MTU is compared to the path MTU associated with the SDP. The path MTU can be administratively defined in the context of the SDP. The default or administrative path MTU can be dynamically reduced due to the MTU capabilities discovered by the tunneling mechanism of the SDP or the egress interface MTU capabilities based on the next hop in the tunnel path. If the service MTU is larger than the path MTU, the SDP binding for the service will be placed in an inoperative state. If the service MTU is equal to or less than the path MTU, then the SDP binding will be placed in an operational state.

In the event that a service MTU, port or channel MTU, or path MTU is dynamically or administratively modified, then all associated SAP and SDP binding operational states are automatically re-evaluated.

The **no** form of this command returns the default **service-mtu** for the indicated service type to the default value.

Default VPLS: 1514

The following table displays MTU values for specific VC types.

VC-Type	Example Service MTU	Advertised MTU
Ethernet	1514	1500
Ethernet (with preserved dot1q)	1518	1504
VPLS	1514	1500
VPLS (with preserved dot1q)	1518	1504
VLAN (dot1p transparent to MTU value)	1514	1500
VLAN (QinQ with preserved bottom Qtag)	1518	1504

octets — The size of the MTU in octets, expressed as a decimal integer.

Values 1 — 9194

split-horizon-group

Syntax [**no**] **split-horizon-group** [*group-name*] [*residential-group*]

Context config>service>vpls

Description This command creates a new split horizon group for the VPLS instance. Traffic arriving on a SAP or spoke SDP within this split horizon group will not be copied to other SAPs or spoke SDPs in the same split horizon group.

A split horizon group must be created before SAPs and spoke SDPs can be assigned to the group.

The split horizon group is defined within the context of a single VPLS. The same group-name can be re-used in different VPLS instances.

Up to 30 split horizon groups can be defined per VPLS instance.

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

Parameters *group-name* — Specifies the name of the split horizon group to which the SDP belongs.

residential-group — Defines a split horizon group as a residential split horizon group (RSHG). Doing so entails that:

a) SAPs which are members of this Residential Split Horizon Group will have:

- Double-pass queuing at ingress as default setting (can be disabled)
- STP disabled (can not be enabled)
- ARP reply agent enabled per default (can be disabled)
- MAC pinning enabled per default (can be disabled)
- Besides the multicast downstream also broadcast packets are discarded thus also blocking the unknown, flooded traffic

b) Spoke SDPs which are members of this Residential Split Horizon Group will have:

- Downstream multicast traffic supported
- Double-pass queuing is not applicable
- STP is disabled (can be enabled)
- ARP reply agent is not applicable (dhcp-lease-states are not supported on spoke SDPs)
- MAC pinning enabled per default (can be disabled)

Default A split horizon group is by default not created as a residential-group.

restrict-protected-src

Syntax **restrict-protected-src** *alarm-only*
no restrict-protected-src

Context config>service>vpls>split-horizon-group
config>service>vpls>sap

Description This command indicates how the agent will handle relearn requests for protected MAC addresses. While enabled all packets entering the configured SAP, or any SAP that is part of the configured split horizon group (SHG) will be verified not to contain a protected source MAC address. If the packet is found to contain such an address the packet will be discarded, a trap will be generated and the SAP will be made operationally down. The SAP must be shutdown and enabled (**no shutdown**) for this state to be cleared. Notice that the use of restrict-protected-src under the SAP is designed only to allow this capability for SAPs that are not part of a SHG. If a SAP is part of a SHG that has **restrict-protected-src** enabled the SAP will be restricted.

Default **no restrict-protected-src**

restrict-unprotected-dst

Syntax **restrict-unprotected-dst** *alarm-only*
no restrict-unprotected-dst

Context config>service>vpls>split-horizon-group
config>service>vpls>sap

Description This command indicates how the system will forward packets destined to an unprotected MAC address. While enabled all packets entering the configured SAP will be verified to contain a protected destination MAC address. If the packet is found to contain a non-protected destination MAC, it will be discarded. Detecting a non-protected destination MAC on the SAP will not cause the SAP to be placed in the operationally down state.

If a MAC is defined with a static address, after that MAC was already learned by the system on a restricted SAP, the system will bring the SAP to an operationally down state, flush the MAC and generate a trap.

If the destination MAC address of a packet entering a restricted SAP with **restrict-unprotected-dst** enabled, and if it is not a protected MAC, the packet will be discarded.

If the destination MAC address is unknown, even if the packet is entering a restricted SAP, with **restrict-unprotected-dst** enabled, it will be flooded.

You can specify a static MAC in the **mac-protect** list. This action fails if the static MAC is defined under a restricted SAP. If the MAC is protected first, the system does not allow it to be added as static under a restricted SAP.

Default **no restrict-unprotected-dst**

tod-suite

Syntax **tod-suite** *tod-suite-name*
no tod-suite

Context config>service>vpls>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>cron** context.

Default **no tod-suite**

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

General Switch Management Protocol Commands

gsmp

Syntax **gsmp**

Context config>service>vpls

Description This command enables the context to configure General Switch Management Protocol (GSMP) connections maintained in this service.

Default **not enabled**

group

Syntax [**no**] **group** *name*

Context config>service>vpls>gsmp

Description This command specifies a GSMP name. A GSMP group name is unique only within the scope of the service in which it is defined.

ancp

Syntax **ancp**

Context config>service>vpls>gsmp>group

Description This command configures Access Node Control Protocol (ANCP) parameters for this GSMP group.

dynamic-topology-discover

Syntax [**no**] **dynamic-topology-discover**

Context config>service>vpls>gsmp>group>ancp

Description This command enables the ANCP dynamic topology discovery capability.

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

line-configuration

Syntax [**no**] **line-configuration**

Context config>service>vpls>gsmp>group>ancp

Description This command enables the ANCP line-configuration capability.

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

oam

Syntax [**no**] **oam**

Context config>service>vpls>gsmp>group>ancp

Description This command specifies whether or not the GSMP ANCP OAM capability should be negotiated at startup of the GSMP connection.

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

hold-multiplier

Syntax **hold-multiplier** *multiplier*
no hold-multiplier

Context config>service>vpls>gsmp>group

Description This command configures the hold-multiplier for the GSMP connections in this group.

Parameters *multiplier* — Specifies the GSMP hold multiplier value.

Values 1 — 100

keepalive

Syntax **keepalive** *seconds*
no keepalive

Context config>service>vpls>gsmp>group

Description This command configures keepalive values for the GSMP connections in this group.

Parameters *seconds* — Specifies the GSMP keepalive timer value in seconds.

Values 1 — 25

neighbor

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

Context config>service>vpls>gsmp>group

Description This command configures a GSMP ANCP neighbor.

Parameters *ip-address* — Specifies the IP address of the GSMP ANCP neighbor.

local-address

Syntax **local-address** *ip-address*
no local-address

Context config>service>vpls>gsmp>group>neighbor

Description This command configures the source ip-address used in the connection towards the neighbor. The local address is optional. If specified the node will accept connections only for that address in the service running ANCP. The address may be created after the reference but connections will not be accepted until it is created. If the local address is not used, the system accepts connections on any interface within the routing context.

Parameters *ip-address* — specifies the source IP address to be used in the connection toward the neighbor.

priority-marking

Syntax **priority-marking dscp** *dscp-name*
priority-marking prec *ip-prec-value*
no priority-marking

Context config>service>vpls>gsmp>group>neighbor

Description This command configures the type of priority marking to be used.

Parameters **dscp** *dscp-name* — Specifies the DSCP code-point to be used.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

prec *ip-prec-value* — Specifies the precedence value to be used.

Values 0 — 7

VPLS DHCP Commands

dhcp

Syntax **dhcp**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description This command enables the context to configure DHCP parameters.

lease-populate

Syntax **lease-populate** [*nmb-of-entries*]
no lease-populate

Context config>service>vpls>sap>dhcp

Description This command enables and disables dynamic host lease state management for VPLS SAPs. For VPLS, DHCP snooping must be explicitly enabled (using the **snoop** command) at all points where DHCP messages requiring snooping enter the VPLS instance (both from the DHCP server and from the subscribers). Lease state information is extracted from snooped DHCP ACK messages to populate lease state table entries for the SAP.

The optional number-of-entries parameter is used to define the number of lease state table entries allowed for this SAP or IP interface. If number-of-entries 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 DHCP ACK messages are discarded.

The retained lease state information representing dynamic hosts may be used to:

- populate a SAP based anti-spoof filter table to provide dynamic anti-spoof filtering. If the system is unable to populate the dynamic host information in the anti-spoof filter table on the SAP, the DHCP ACK message must be discarded without adding a new lease state entry or updating an existing lease state entry.
- generate dynamic ARP replies if **arp-reply-agent** is enabled.

Default **no lease-populate**

Parameters *nbr-of-entries* — Specifies the number of DHCP leases allowed.

Values 1 — 8000

option

Syntax [**no**] **option**

Context config>service>vpls>sap>dhcp

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

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

Default **no option**

action

Syntax **action** [*dhcp-action*]
no action

Context config>service>vpls>sap>dhcp>option

Description This command configures the Relay Agent Information Option (Option 82) processing.

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

Default The default is to keep the existing information intact.

Parameters *dhcp-action* — Specifies the DHCP option action.

replace — In the upstream direction (from the user), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.

keep — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is sent on towards the client.

The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.

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

circuit-id

Syntax **circuit-id** [**ascii-tuple** | **vlan-ascii-tuple**]

Context config>service>vpls>sap>dhcp>option

Description When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** suboption of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by

"|". If no keyword is configured, then the circuit-id suboption will not be part of the information option (Option 82).

If disabled, the **circuit-id** suboption of the DHCP packet will be left empty.

Default **no circuit-id**

Parameters **ascii-tuple —** Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used.

vlan-ascii-tuple — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and and qinq encapsulated ports only. Thus, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

remote-id

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

Context config>service>vpls>sap>dhcp>option

Description This command specifies what information goes into the remote-id suboption in the DHCP Relay packet.

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

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

Default **no remote-id**

Parameters **mac —** This keyword specifies the MAC address of the remote end is encoded in the suboption.

string *string* **—** Specifies the remote-id.

vendor-specific-option

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

Context config>service>vpls>sap>dhcp>option
config>service>ies>if>dhcp>option

Description This command configures the vendor specific suboption of the DHCP relay packet.

client-mac-address

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

Context config>service>vpls>sap>dhcp>option>vendor

Description This command enables the sending of the MAC address in the vendor specific suboption of the DHCP relay packet.

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

sap-id

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

Context config>service>vpls>sap>dhcp>option>vendor

Description This command enables the sending of the SAP ID in the vendor specific suboption of the DHCP relay packet.

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

service-id

Syntax [**no**] **service-id**

Context config>service>vpls>sap>dhcp>option>vendor

Description This command enables the sending of the service ID in the vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the vendor specific suboption of the DHCP relay packet.

string

Syntax [**no**] **string** *text*

Context config>service>vpls>sap>dhcp>option>vendor

Description This command specifies the string in the vendor specific suboption of the DHCP relay packet.

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

Parameters *text* — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

system-id

Syntax [**no**] **system-id**

Context config>service>vpls>sap>dhcp>option>vendor

Description This command specifies whether the system-id is encoded in the vendor specific sub-option of Option 82.

proxy-server

Syntax **proxy-server**

Context config>service>vpls>sap>dhcp

Description This command configures the DHCP proxy server.

emulated-server

Syntax **emulated-server** *ip-address*
no emulated-server

Context config>service>vpls>sap>dhcp>proxy

Description This command configures the IP address which will be used as the DHCP server address in the context of this VPLS SAP. Typically, the configured address should be in the context of the subnet represented by the VPLS.

The **no** form of of this command reverts to the default setting. The local proxy server will not become operational without the emulated-server address being specified.

Parameters *ip-address* — Specifies the emulated server address.

lease-time

Syntax **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*] [**radius-override**]
no lease-time

Context config>service>vpls>sap>dhcp>proxy

Description This command defines the length of lease time that will be provided to DHCP clients. By default, the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.

The **no** form of this command disables the use of the lease-time command. The local proxy server will use the lease time offered by either a RADIUS or DHCP server.

Default 7 days 0 hours 0 seconds

Parameters *days* — Specifies the number of days that the given IP address is valid.

Values 0 — 3650

hours — Specifies the number of hours that the given IP address is valid.

Values 0 — 23

minutes — Specifies the number of minutes that the given IP address is valid.

Values 0 — 59

seconds — Specifies the number of seconds that the given IP address is valid.

Values 0 — 59

snoop

Syntax [**no**] **snoop**

Context config>service>vpls>sap>dhcp
config>service>vpls>spoke-sdp>dhcp
config>service>vpls>mesh-sdp>dhcp

Description This command enables DHCP snooping of DHCP messages on the SAP or SDP. Enabling DHCP snooping on VPLS interfaces (SAPs and SDP bindings) is required where DHCP messages important to lease state table population are received, or where Option 82 information is to be inserted. This includes interfaces that are in the path to receive messages from either DHCP servers or from subscribers.

Use the **no** form of the command to disable DHCP snooping on the specified VPLS SAP or SDP binding.

Default **no snoop**

VPLS STP Commands

stp

Syntax **stp**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command enables the context to configure the Spanning Tree Protocol (STP) parameters within a VPLS spoke service distribution point (SDP). Alcatel-Lucent's STP is simply the Spanning Tree Protocol (STP) with a few modifications to better suit the operational characteristics of VPLS services. The most evident change is to the root bridge election. Since the core network operating between Alcatel-Lucent's service routers should not be blocked, the root path is calculated from the core perspective.

auto-edge

Syntax **auto-edge**
no auto-edge

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command configures automatic detection of the edge port characteristics of the SAP or spoke SDP.

If auto-edge is enabled, and STP concludes there is no bridge behind the spoke SDP, the OPER_EDGE variable will dynamically be set to true. If auto-edge is enabled, and a BPDU is received, the OPER_EDGE variable will dynamically be set to true (see edge-port on page 483).

The **no** form of this command returns the auto-detection setting to the default value.

Default **auto-edge**

edge-port

Syntax [**no**] **edge-port**

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command configures the SAP or SDP as an edge or non-edge port. If **auto-edge** is enabled for the SAP, this value will be used only as the initial value.

NOTE: The function of the **edge-port** command is similar to the **rapid-start** command. It tells RSTP that it is on the edge of the network (for example, there are no other bridges connected to that port)

and, as a consequence, it can immediately transition to a forwarding state if the port becomes available.

RSTP, however, can detect that the actual situation is different from what **edge-port** may indicate.

Initially, the value of SAP parameter is set to edge-port. This value will change if:

- A BPDU is received on that port. This means that after all there is another bridge connected to this port. Then the edge-port becomes disabled.
- If auto-edge is configured and no BPDU is received within a certain period of time, RSTP concludes that it is on an edge and enables the SAP edge-port.

The **no** form of this command returns the edge port setting to the default value.

Default **no edge-port**

forward-delay

Syntax **forward-delay** *seconds*
no forward-delay

Context config>service>vpls>stp

Description RSTP, as defined in the IEEE 802.1D-2004 standards, will normally transition to the forwarding state via a handshaking mechanism (rapid transition), without any waiting times. If handshaking fails (e.g. on shared links, see below), the system falls back to the timer-based mechanism defined in the original STP (802.1D-1998) standard.

A shared link is a link with more than two nodes (for example, a shared 10/100BaseT segment). The `port-type` command is used to configure a link as point-to-point or shared.

For timer-based transitions, the 802.1D-2004 standard defines an internal variable forward-delay, which is used in calculating the default number of seconds that a SAP or spoke SDP spends in the discarding and learning states when transitioning to the forwarding state.

The value of the forward-delay variable depends on the STP operating mode of the VPLS instance:

- in `rstp` or `mstp` mode, but only when the SAP or spoke SDP has not fallen back to legacy STP operation, the value configured by the `hello-time` command is used;
- in all other situations, the value configured by the `forward-delay` command is used.

Default 15 seconds

Parameters *seconds* — The forward delay timer for the STP instance in seconds. Allowed values are integers in the range of 4 to 30 .

hello-time

Syntax **hello-time** *hello-time*
no hello-time

Context config>service>vpls>stp

Description This command configures the Spanning Tree Protocol (STP) hello time for the Virtual Private LAN Service (VPLS) STP instance.

The hello time parameter defines the default timer value that controls the sending interval between BPDU configuration messages by this bridge, on ports where this bridge assumes the designated role.

The active hello time for the spanning tree is determined by the root bridge (except when the STP is running in RSTP mode, then the hello time is always taken from the locally configured parameter).

The configured hello-time can also be used to calculate the forward delay. See auto-edge on page 483.

The **no** form of this command returns the hello time to the default value.

Default **2 seconds**

Parameters *hello-time* — The hello time for the STP instance in seconds.

Values 1 — 10

hold-count

Syntax **hold-count** *BDPU tx hold count*
no hold-count

Context config>service>vpls>stp

Description This command configures the peak number of BPDUs that can be transmitted in a period of one second.

The **no** form of this command returns the hold count to the default value

Default **6**

Parameters *BDPU tx hold count* — The hold count for the STP instance in seconds.

Values 1 — 10

link-type

Syntax **link-type** {**pt-pt** | **shared**}
no link-type

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command instructs STP on the maximum number of bridges behind this SAP or spoke SDP. If there is only a single bridge, transitioning to forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected via a shared media, their SAP or spoke SDPs should all be configured as shared, and timer-based transitions are used.

The **no** form of this command returns the link type to the default value.

Default **pt-pt**

mst-instance

Syntax **mst-instance** *mst-inst-number*

Context config>service>vpls>sap>stp

Description This command enables the context to configure MSTI related parameters at SAP level. This context can be open only for existing mst-instances defined at the service level (see **mst-instance**).

Default **none**

Parameters *mst-inst-number* — Specifies an existing Multiple Spanning Tree Instance number.

Values 1 — 4094

mst-path-cost

Syntax **mst-path-cost** *inst-path-cost*
no mst-path-cost

Context config>service>vpls>sap>stp>mst-instance

Description This commands specifies path-cost within a given instance, expressing probability that a given port will be put into the forwarding state in case a loop occurs (the highest value expresses lowest priority).

The **no** form of this command sets port-priority to its default value.

Default The path-cost is proportional to link speed.

Parameters *inst-path-cost* — Specifies the contribution of this port to the MSTI path cost of paths towards the spanning tree regional root which include this port.

Values 1 — 200000000

mst-port-priority

Syntax **mst-port-priority** *stp-priority*
no mst-port-priority

Context config>service>vpls>sap>stp>mst-instance

Description This commands specifies the port priority within a given instance, expressing probability that a given port will be put into the forwarding state if a loop occurs.

The **no** form of this command sets port-priority to its default value.

Default **128**

Parameters *stp-priority* — Specifies the value of the port priority field.

max-age

Syntax **max-age** *seconds*
no max-age

Context config>service>vpls>stp

Description This command indicates how many hops a BPDU can traverse the network starting from the root bridge. The message age field in a BPDU transmitted by the root bridge is initialized to 0. Each other bridge will take the message_age value from BPDUs received on their root port and increment this value by 1. The message_age thus reflects the distance from the root bridge. BPDUs with a message age exceeding max-age are ignored.

STP uses the max-age value configured in the root bridge. This value is propagated to the other bridges via the BPDUs.

The **no** form of this command returns the max ageto the default value.

Default 20 seconds

Parameters *seconds* — The max info age for the STP instance in seconds. Allowed values are integers in the range 6 to 40.

mode

Syntax **mode** {**rstp** | **comp-dot1w** | **dot1w** | **mstp**}
no mode

Context config>service>vpls>stp

Description This command specifies the version of Spanning Tree Protocol the bridge is currently running.

See section Spanning Tree Operating Modes on page 335 for details on these modes.

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

Default **rstp**

Parameters **rstp —** Corresponds to the Rapid Spanning Tree Protocol specified in IEEE 802.1D/D4-2003.

dot1w — Corresponds to the mode where the Rapid Spanning Tree is backward compatible with IEEE 802.1w.

compdot1w — Corresponds to the Rapid Spanning Tree Protocol fully conformant to IEEE 802.1w.

mstp — Sets MSTP as the STP mode of operation. Corresponds to the Multiple Spanning Tree Protocol specified in 802.1Q REV/D5.0-09/2005.

mst-instance

Syntax [**no**] **mst-instance** *mst-inst-number*

Context config>service>vpls>stp

Description This command creates the context to configure MST instance (MSTI) related parameters. Up to 16 instances will be supported by MSTP. The instance 0 is mandatory by protocol and therefore, it cannot be created by the CLI. The software will maintain this instance automatically.

Default **none**

Parameters *mst-inst-number* — Specifies the Multiple Spanning Tree instance.

Values 1 — 4094

mst-priority

Syntax **mst-priority** *bridge-priority*
no mst-priority

Context config>service>vpls>stp>mst-instance

Description This command specifies the bridge priority for this specific Multiple Spanning Tree Instance for this service. The *bridge-priority* value reflects likelihood that the switch will be chosen as the regional root switch (65535 represents the least likely). It is used as the highest 4 bits of the Bridge ID included in the MSTP BPDU's generated by this bridge.

The priority can only take on values that are multiples of 4096 (4k). If a value is specified that is not a multiple of 4K, then the value will be replaced by the closest multiple of 4K, which is lower than the value entered.

The **no** form of this command sets the bridge-priority to its default value.

Default 32768 — All instances created by **vlan-range** command and not having explicit definition of bridge-priority will inherit default value.

Parameters *bridge-priority* — Specifies the priority of this specific Multiple Spanning Tree Instance for this service.

Values 0 — 65535

vlan-range

Syntax [**no**] **vlan-range** [*vlan-range*]

Context config>service>vpls>stp>mst-instance

Description This command specifies a range of VLANs associated with a certain MST-instance. This range applies to all SAPs of the mVPLS.

Every VLAN range that is not assigned within any of the created **mst-instance** is automatically assigned to mst-instance 0. This instance is automatically maintained by the software and cannot be modified. Changing the VLAN range value can be performed only when the given mst-instance is shutdown.

The **no** form of this command removes the **vlan-range** from given **mst-instance**.

Parameters *vlan-range* — The first VLAN range specifies the left-bound (i.e., minimum value) of a range of VLANs that are associated with the mVPLS SAP. This value must be smaller than (or equal to) the second VLAN range value.

The second VLAN range specifies the right-bound (i.e., maximum value) of a range of VLANs that are associated with the mVPLS SAP.

Values 1 to 4094 — 1 to 4094

mst-max-hops

Syntax **mst-max-hops** *hops-count*
no mst-max-hops

Context config>service>vpls>stp

Description This command specifies the number of hops in the region before BPDU is discarded and the information held for the port is aged out. The root bridge of the instance sends a BPDU (or M-record) with remaining-hop-count set to configured <*max-hops*>. When a bridge receives the BPDU (or M-record), it decrements the received remaining-hop-count by 1 and propagates it in BPDU (or M-record) it generates.

The **no** form of this command sets the *hops-count* to its default value.

Default **20**

Parameters *hops-count* — Specifies the maximum number of hops.

Values 1 — 40

mst-name

Syntax **mst-name** *region-name*
no mst-name

Context config>service>vpls>stp

Description This command defines an MST region name. Two bridges are considered as a part of the same MST region as soon as their configuration of the MST region name, the MST-revision and VLAN-to-instance assignment is identical.

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

Default **no mst-name**

Parameters *region-name* — Specifies an MST-region name up to 32 characters in length.

mst-revision

Syntax **mst-revision** *revision-number*

Context config>service>vpls>stp

Description This command defines the MST configuration revision number. Two bridges are considered as a part of the same MST region as soon as their configuration of MST-region name, MST-revision and VLAN-to-instance assignment is identical.

The **no** form of this command returns MST configuration revision to its default value.

Default **0**

Parameters *revision-number* — Specifies the MSTP region revision number to define the MSTP region.

Values 0 — 65535

path-cost

Syntax **path-cost** *sap-path-cost*
no path-cost

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command configures the Spanning Tree Protocol (STP) path cost for the SAP or spoke SDP.

The SAP or spoke SDP path cost is used by STP to calculate the path cost to the root bridge. The path cost in BPDUs received on the root port is incremented with the configured path cost for that SAP or spoke SDP. When BPDUs are sent out other egress SAPs or spoke SDPs, the newly calculated root path cost is used. These are the values used for CIST when running MSTP.

STP suggests that the path cost is defined as a function of the link bandwidth. Since SAPs and spoke SDPs are controlled by complex queuing dynamics, in the SR-Series the STP path cost is a purely static configuration.

The **no** form of this command returns the path cost to the default value.

path-cost — The path cost for the SAP or spoke SDP.

Values 1 — 65535 (1 is the lowest cost)

Default 10

port-num

Syntax [**no**] **port-num** *virtual-port-number*

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command configures the virtual port number which uniquely identifies a SAP within configuration bridge protocol data units (BPDUs). The internal representation of a SAP is unique to a system and has a reference space much bigger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a SAP and identifies it with it's own virtual port number that is unique to every other SAP defined on the TLS. The virtual port number is assigned at the time that the SAP is added to the TLS. Since the order that the SAP was added to the TLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance.

The virtual port number cannot be administratively modifed.

priority

Syntax **priority** *bridge-priority*
no priority

Context config>service>vpls>stp

Description The bridge-priority command is used to populate the priority portion of the bridge ID field within outbound BPDUs (the most significant 4 bits of the bridge ID). It is also used as part of the decision process when determining the best BPDU between messages received and sent.

All values will be truncated to multiples of 4096, conforming with IEEE 802.1t and 802.1D-2004.

The **no** form of this command returns the bridge priority to the default value.

Default By default, the bridge priority is configured to 4096 which is the highest priority.

Parameters *bridge-priority* — The bridge priority for the STP instance.

Values Allowed values are integers in the range of 4096 — 65535 with 4096 being the highest priority. The actual bridge priority value stored/used is the number entered with the lowest 12 bits masked off which means the actual range of values is 4096 to 61440 in increments of 4096.

priority

Syntax **priority** *stp-priority*
no priority

Context config>service>vpls>spoke-sdp
config>service>vpls>sap>stp

Description This command configures the Alcatel-Lucent Spanning Tree Protocol (STP) priority for the SAP or spoke SDP.

STP priority is a configurable parameter associated with a SAP or spoke SDP. When configuration BPDUs are received, the priority is used in some circumstances as a tie breaking mechanism to determine whether the SAP or spoke SDP will be designated or blocked.

In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (0 to 255 also) to create a 16 bit value. In the latest STP standard (802.1D-2004) only the upper 4 bits of the port priority field are used to encode the SAP or Spoke SDP priority. The remaining 4 bits are used to extend the port ID field into a 12 bit virtual port number field. The virtual port number uniquely references a SAP or spoke SDP within the STP instance.

STP computes the actual priority by taking the input value and masking out the lower four bits.The result is the value that is stored in the SDP priority parameter. For instance, if a value of 0 is entered, masking out the lower 4 bits results in a parameter value of 0. If a value of 255 is entered, the result is 240.

The **no** form of this command returns the STP priority to the default value.

Default **128**

Parameters *stp-priority* — The STP priority value for the SAP or Spoke SDP. Allowed values are integer in the range of 0 to 255, 0 being the highest priority. The actual value used for STP priority (and stored

in the configuration) will be the result of masking out the lower 4 bits - thus the actual value range is 0 to 240 in increments of 16.

Default 128

VPLS SAP Commands

sap

Syntax **sap** *sap-id* [**split-horizon-group** *group-name*] [**create**]
no sap *sap-id*

Context config>service>vpls

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the 7750. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface** *port-type port-id* **mode access** command. Channelized TDM ports are always access ports.

If a port is 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.

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 Ethernet Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed.

Default No SAPs are defined.

Special Cases A VPLS SAP can be defined with Ethernet ports, SONET/SDH or TDM channels. Up to 127 SAPs can be defined in a VPLS service per Media Dependent Adapter (MDA). Attempts to create more than 127 SAPs will generate an error. Up to 49 SDPs can be associated with a VPLS in a single 7750 SR router. Each SDP must have a unique 7750 SR destination or an error will be generated. Split horizon groups can only be created in the scope of a VPLS service.

A default SAP has the following format: port-id:*. This type of SAP is supported only on Ethernet MDAs and its creation is allowed only in the scope of Layer 2 services (Epipe and VPLS). This type of SAP is mutually exclusive with a SAP defined by explicit null encapsulation (for example, 1/1/1:0).

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	*slot*/*mda*/*port*[.*channel*]	2/2/11 1/2/3.1
null	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id* | *aps-id*]	*port-id*: 1/2/3 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*: lag-100 *aps-id*: aps-1
dot1q	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id* | *aps-id*]:qtag1	*port-id*:qtag1: 1/2/3:100 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*:qtag1:lag-100:102 *aps-id*:qtag1: aps-1:103
qinq	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	*port-id*:qtag1.qtag2: 1/2/3:100.10 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*:qtag1.qtag2:lag-100:
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:vpi/vci |vpi |vpi1.vpi2]	port-id: 9/1/1 aps-id: aps-1 *bundle-id*: bundle-ima-1/1.1 bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 vpi/vci: 16/26 vpi: 16 vpi1.vpi2: 16.200
frame-relay	[*port-id* | *aps-id*]:*dlci*	*port-id*: 1/1/1:100 *aps-id*: aps-1 *dlci*: 16
cisco-hdlc	*slot/mda/port.channel*	*port-id*: 1/2/3.1

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	

	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

bundle-id — Specifies the multilink bundle to be associated with this IP interface. The **bundle** keyword must be entered at the beginning of the parameter.
The command syntax must be configured as follows:

bundle-id: **bundle**-*type-slot-id*/*mda-slot.bundle-num*
bundle-id value range: 1 — 128

For example:

```
*A:ALA-12>config#  port bundle-ppp-5/1.1
*A:ALA-12>config>port# multilink-bundle
```

bgprp-id — Specifies the bundle protection group ID to be associated with this IP interface. The **bpgrp** keyword must be entered at the beginning of the parameter.

The command syntax must be configured as follows:

bpgrp-id: **bpgrp**-*type-bpgrp-num*
type: ima
bpgrp-num value range: 1 — 1280

For example:

```
*A:ALA-12>config# port bpgrp-ima-1
*A:ALA-12>config>service>vpls$ sap bpgrp-ima-1
```

qtag1, qtag2 — Specifies the encapsulation value used to identify the SAP on the port or sub-port. If this parameter is not specificially defined, the default value is 0.

Values qtag1: 0 — 4094
qtag2 : * | 0 — 4094

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 TDM	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 TDM	BCP-Dot1q	0 — 4094	The SAP is identified by the 802.1Q tag on the channel.
SONET/SDH TDM	Frame Relay	16 — 991	The SAP is identified by the data link connection identifier (DLCI).
SONET/SDH ATM	ATM	vpi (NNI) 0 — 4095 vpi (UNI) 0 — 255 vci 1, 2, 5 — 65535 -	The SAP is identified by port or by PVPC or PVCC identifier (vpi, vpi/vci, or vpi range)

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

split-horizon-group *group-name* **—** Specifies the name of the split horizon group to which the SAP belongs.

discard-unknown-source

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

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description When this command is enabled, packets received on a SAP or or a spoke SDP with an unknown source MAC address will be dropped only if the maximum number of MAC addresses for that SAP or spoke SDP (see max-nbr-mac-addr on page 498) has been reached. If max-nbr-mac-addr has not been set for the SAP or spoke SDP, enabling discard-unknown-source has no effect.

When disabled, the packets are forwarded based on the destination MAC addresses.

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

Default **no discard-unknown**

limit-mac-move

Syntax **limit-mac-move** [**blockable** | **non-blockable**]
no limit-mac-move

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command indicates whether or not the mac-move agent, when enabled using **config>service>vpls>mac-move** or **config>service>epipe>mac-move**, will limit the MAC re-learn (move) rate on this SAP.

Default **SAPs and spoke SDPs are blockable**

Parameters **blockable —** The agent will monitor the MAC re-learn rate on the SAP, and it will block it when the re-learn rate is exceeded.

non-blockable — When specified, this SAP will not be blocked, and another blockable SAP will be blocked instead.

mac-pinning

Syntax [**no**] **mac-pinning**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description Enabling this command will disable re-learning of MAC addresses on other SAPs within the VPLS. The MAC address will remain attached to a given SAP for duration of its age-timer.

The age of the MAC address entry in the FIB is set by the age timer. If **mac-aging** is disabled on a given VPLS service, any MAC address learned on a SAP/SDP with **mac-pinning** enabled will remain in the FIB on this SAP/SDP forever.

Every event that would otherwise result in re-learning will be logged (MAC address; original-SAP; new-SAP).

Note that MAC addresses learned during DHCP address assignment (DHCP snooping enabled) are not impacted by this command. MAC-pinning for such addresses is implicit.

Default When a SAP or spoke SDP is part of a Residential Split Horizon Group (RSHG), MAC pinning is activated at creation of the SAP. Otherwise MAC pinning is not enabled by default.

max-nbr-mac-addr

Syntax **max-nbr-mac-addr** *table-size*
no max-nbr-mac-addr

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this SAP or spoke SDP.

When the configured limit has been reached, and discard-unknown-source has been enabled for this SAP or spoke SDP (see discard-unknown-source on page 496), packets with unknown source MAC addresses will be discarded.

The **no** form of the command restores the global MAC learning limitations for the SAP or spoke SDP.

Default **no max-nbr-mac-addr**

Parameters *table-size* — Specifies the maximum number of learned and static entries allowed in the FDB of this service.

Values 1 — 196607

Chassis-mode C limit: 196607

multi-service-site

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

Context config>service>vpls>sap

Description This command associates the SAP with a *customer-site-name*. If the specified *customer-site-name* does not exist in the context of the service customer ID an error occurs and the command will not execute. If *customer-site-name* exists, the current and future defined queues on the SAP (ingress and egress) will attempt to use the scheduler hierarchies created within *customer-site-name* as parent schedulers.

This command is mutually exclusive with the SAP ingress and egress **scheduler-policy** commands. If a **scheduler-policy** has been applied to either the ingress or egress nodes on the SAP, the **multi-service-site** command will fail without executing. The locally applied scheduler policies must be removed prior to executing the **multi-service-site** command.

The **no** form of the command removes the SAP from any multi-service customer site the SAP belongs to. Removing the site can cause existing or future queues to enter an orphaned state.

Default **None**

customer-site-name — The customer-site-name must exist in the context of the customer-id defined as the service owner. If customer-site-name exists and local scheduler policies have not been applied to the SAP, the current and future queues defined on the SAP will look for their parent schedulers within the scheduler hierarchies defined on customer-site-name.

Values Any valid customer-site-name created within the context of the customer-id

static-mac

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

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command creates a local static MAC entry in the Virtual Private LAN Service (VPLS) forwarding database (FDB) associated with the Service Access Point (SAP).

In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Local static MAC entries create a permanent MAC address to SAP association in the forwarding database for the VPLS instance so that MAC address will not be learned on the edge device.

Note that static MAC defnitions on one edge device are not propagated to other edge devices participating in the VPLS instance, that is, each edge device has an independent forwarding database for the VPLS.

Only one static MAC entry (local or remote) can be defined per MAC address per VPLS instance.

By default, no static MAC address entries are defined for the SAP.

The **no** form of this command deletes the static MAC entry with the specified MAC address associated with the SAP from the VPLS forwarding database.

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

managed-vlan-list

Syntax **managed-vlan-list**

Context config>service>vpls>sap

Description This command enables the context to configure VLAN ranges to be managed by a management VPLS. The list indicates, for each SAP, the ranges of associated VLANs that will be affected when the SAP changes state. This managed-vlan-list is not used when STP mode is MSTP in which case the vlan-range is taken from the **config>service>vpls>stp>msti** configuration.

This command is only valid when the VPLS in which it is entered was created as a management VPLS.

default-sap

Syntax [**no**] **default-sap**

Context config>service>vpls>sap>managed-vlan-list

Description This command adds a default SAP to the managed VLAN list.

The no form of the command removes the default SAP to the managed VLAN list.

range

Syntax [**no**] **range** *vlan-range*

Context config>service>vpls>sap>managed-vlan-list

Description This command configures a range of VLANs on an access port that are to be managed by an existing management VPLS.

This command is only valid when the VPLS in which it is entered was created as a management VPLS, and when the SAP in which it was entered was created on an Ethernet port with encapsulation type of dot1q or qinq, or on a Sonet/SDH port with encapsulation type of bcp-dot1q.

To modify the range of VLANs, first the new range should be entered and afterwards the old range removed. See Modifying VPLS Service Parameters on page 436.

Default **None**

Parameters *vlan-range* — Specify the VLAN start value and VLAN end value. The end-vlan must be greater than start-vlan. The format is <start-vlan>-<end-vlan>

Values start-vlan: 0 — 4094
end-vlan: 0 — 4094

VPLS SAP ATM Commands

atm

Syntax **atm**

Context config>service>vpls>sap

Description This command enables access to the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as:

- Configuring ATM port or ATM port-related functionality on MDAs supporting ATM functionality
- Configuring ATM-related configuration for ATM-based SAPs that exist on MDAs supporting ATM functionality.

If ATM functionality is not supported for a given context, the command returns an error.

egress

Syntax **egress**

Context config>service>vpls>sap>atm

Description This command enables the context to configure egress ATM attributes for the SAP.

encapsulation

Syntax **encapsulation** *atm-encap-type*

Context config>service>vpls>sap>atm

Description This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition references RFC 2684, *Multiprotocol Encapsulation over ATM AAL5*, and to the ATM Forum LAN Emulation specification.

Ingress traffic that does not match the configured encapsulation will be dropped.

Default The encapsulation is driven by the services for which the SAP is configured.
For IES and VPRN service SAPs, the default is **aal5snap-routed**.

Parameters *atm-encap-type* — Specify the encapsulation type.

Values **aal5snap-routed** — Routed encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684.
aal5mux-ip — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684.

ingress

Syntax **ingress**

Context config>service>vpls>sap>atm

Description This command enables the context to configure ingress ATM attributes for the SAP.

traffic-desc

Syntax **traffic-desc** *traffic-desc-profile-id*
no traffic-desc

Context config>service>vpls>sap>atm>ingress
config>service>vpls>sap>atm>egress

Description This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP).

When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction.

When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.

The **no** form of the command reverts the traffic descriptor to the default traffic descriptor profile.

Default The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.

Parameters *traffic-desc-profile-id* — Specify a defined traffic descriptor profile (see the QoS atm-td-profile command).

oam

Syntax **oam**

Context config>service>vpls>sap>atm

Description This command enables the context to configure OAM functionality for a PVCC delimiting a SAP.

The ATM-capable MDAs support F5 end-to-end OAM functionality (AIS, RDI, Loopback):

- ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95
- GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996
- GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

alarm-cells

Syntax [**no**] **alarm-cells**

Context config>service>vpls>sap>atm

Description This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC termination to monitor and report the status of their connection by propagating fault information through the network and by driving PVCC's operational status.

When alarm-cells functionality is enabled, a PVCC's operational status is affected when a PVCC goes into an AIS or RDI state because of an AIS/RDI processing (assuming nothing else affects PVCC's operational status, for example, if the PVCC goes DOWN, or enters a fault state and comes back UP, or exits that fault state). RDI cells are generated when PVCC is operationally DOWN. No OAM-specific SNMP trap is raised whenever an endpoint enters/exits an AIS or RDI state, however, if as result of an OAM state change, the PVCC changes operational status, then a trap is expected from an entity the PVCC is associated with (for example a SAP).

The **no** command disables **alarm-cells** functionality for a PVCC. When **alarm-cells** functionality is disabled, the PVCC's operational status is no longer affected by the PVCC's OAM state changes due to AIS/RDI processing. Note that when **alarm-cells** is disabled, a PVCC will change operational status to UP from DOWN due to alarm-cell processing). RDI cells are not generated as result of PVCC going into an AIS or RDI state, however, the PVCC's OAM status will record OAM faults as described above.

Default Enabled for PVCCs delimiting VPLS SAPs.

VPLS Filter and QoS Policy Commands

egress

Syntax **egress**

Context config>service>vpls>sap

Description This command enables the context to configure egress SAP Quality of Service (QoS) policies and filter policies.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing. If no egress filter is defined, no filtering is performed.

ingress

Syntax **ingress**

Context config>service>vpls>sap

Description This command enables the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

agg-rate-limit

Syntax **agg-rate-limit** *agg-rate*
no agg-rate-limit

Context config>service>vpls>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

filter

Syntax

filter ip *ip-filter-id*
filter ipv6 *ipv6-filter-id*
filter mac *mac-filter-id*
no filter [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress
config>service>vpls>mesh-sdp>egress
config>service>vpls>mesh-sdp>ingress
config>service>vpls>spoke-sdp>egress
config>service>vpls>spoke-sdp>ingress

Description This command associates an IP filter policy or MAC 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 or MAC matching criteria. There are two types of filter policies: IP and MAC. Only one type may be applied to a SAP at a time.

The **filter** command is used to associate a filter policy with a specified filter ID with an ingress or egress SAP. The filter ID must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation 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**.

Special Cases **VPLS —** Both MAC and IP filters are supported on a VPLS service SAP.

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

Values 1 — 65535

ipv6 *ipv6-filter-id* **—** Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 — 65535

mac *mac-filter-id* **—** Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1 — 65535

multicast-group

Syntax **multicast-group** *group-name*
no multicast-group

Context config>service>vpls>sap>egress

Description This command places a VPLS Ethernet SAP into an egress multicast group. The SAP must comply with the egress multicast group's common requirements for member SAPs. If the SAP does not comply, the command will fail and the SAP will not be a member of the group. Common requirements for an egress multicast group are listed below:

- If an egress-filter is specified on the egress multicast group, the SAP must have the same egress filter applied.
- If an egress-filter is not defined on the egress multicast group, the SAP cannot have an egress filter applied.
- If the egress multicast group has an encap-type set to null, the SAP must be defined on a port with the port encapsulation type set to null.
- If the egress multicast group has an encap-type set to dot1q, the SAP must be defined on a port with the port encapsulation type set to dot1q and the port's dot1q-etype must match the dot1q-etype defined on the egress multicast group.
- The access port the SAP is created on cannot currently be an egress mirror source.

Once a SAP is a member of an egress multicast group, the following rules apply:

- The egress filter defined on the SAP cannot be removed or modified. Egress filtering is managed at the egress multicast group for member SAPs.
- If the encapsulation type for the access port the SAP is created on is set to dot1q, the port's dot1q-etype value cannot be changed.
- Attempting to define an access port with a SAP that is currently defined in an egress multicast group as an egress mirror source will fail.

Once a SAP is included in an egress multicast group, it is then eligible for efficient multicast replication if the egress forwarding plane performing replication for the SAP is capable. If the SAP is defined as a Link Aggregation Group (LAG) SAP, it is possible that some links in the LAG are on forwarding planes that support efficient multicast replication while others are not. The fact that some or all the forwarding planes associated with the SAP cannot perform efficient multicast replication does not affect the ability to place the SAP into an Egress multicast group.

A SAP may be a member of one and only one egress multicast group. If the multicast-group command is executed with another egress multicast group name, the system will attempt to move the SAP to the specified group. If the SAP is not placed into the new group, the SAP will remain a member of the previous egress multicast group. Moving a SAP into an egress multicast group may cause a momentary gap in replications to the SAP destination while the move is being processed.

The **no** form of the command removes the SAP from any egress multicast group in which it may currently have membership. The SAP will be removed from all efficient multicast replication chains and normal replication will apply to the SAP. A momentary gap in replications to the SAP destination while it is being moved is possible. If the SAP is not currently a member in an egress multicast group, the command has no effect.

Default **no multicast-group**

Parameters *group-name* — The *group-name* is required when specifying egress multicast group membership on a SAP. An egress multicast group with the specified egress-multicast-group-name must exist and the SAP must pass all common requirements or the command will fail.

Values Any valid egress multicast group name.

Default None, an egress multicast group name must be explicitly specified.

qinq-mark-top-only

Syntax [**no**] **qinq-mark-top-only**

Context config>service>vpls>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 to mark during packet egress. When disabled, both set of P-bits are marked. When enabled, only the P-bits in the top Q-tag are marked.

Default **no qinq-mark-top-only**

qos

Syntax **qos** *policy-id* [**shared-queuing** | **multipoint-shared**]
no qos

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

Description This command associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP) or IP interface.

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the *policy-id* does not exist, an error will be returned.

The **qos** command is used to associate both ingress and egress QoS policies. The **qos** command only allows ingress policies to be associated on SAP or IP interface ingress and egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a 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.

When an egress QoS policy is associated with an IES IP interface that has been bound to a VPLS, the policy becomes associated with every SAP on the VPLS and augments the egress 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 or IP interface for ingress or egress, so the default QoS policy is used.

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

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 — This keyword can only be specified on SAP ingress. Specify the ingress shared queue policy used by this SAP. When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

multipoint-shared — This keyword can only be specified on SAP ingress. Multipoint shared queuing is a superset of shared queuing. When multipoint shared queuing keyword is set, in addition to the unicast packets, multipoint packets also used shared queues.

Ingress unicast service queues are mapped one-for-one with hardware queues and unicast packets traverse the ingress forwarding plane twice, similar to the shared-queuing option. In addition, the multipoint queues defined in the ingress SAP QoS policy are not created. Instead, multipoint packets (broadcast, multicast and unknown unicast destined) are treated to the same dual pass ingress forwarding plane processing as unicast packets.

When the value of this object is null, the SAP will use individual ingress QoS queues, instead of the shared ones.

Values Multipoint or not present.

Default Present (the queue is created as non-multipoint).

queue-override

Syntax [**no**] **queue-override**

Context config>service>vpls>sap>egress
config>service>vpls>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** *queue-id*

Context config>service>vpls>sap>egress>queue-override
config>service>vpls>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>vpls>sap>egress>queue-override>queue
config>service>vpls>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** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir — The **cir** parameter defines the constraints enforced when adapting the CIR rate defined within the **queue** *queue-id* **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule — Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset.

Values **max** — The **max** (maximum) keyword is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) keyword is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

avg-frame-overhead

Syntax **avg-frame-overhead** *percent*
no avg-frame-overhead

Context config>service>vpls>sap>egress>queue-override>queue

Description This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a Sonet or SDH port or channel. Queues operating on egress Ethernet

ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- Offered-load — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- Frame encapsulation overhead — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be 10000 x 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 x 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 x 1.1 or 550 octets.
- Frame based within-cir offered-load — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

  As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.
- Frame based PIR — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be 7500 x 1.1 or 8250 octets.
- Frame based within-pir offered-load — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based

offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

Port scheduler operation using frame transformed rates — The port scheduler uses the frame based rates to calculate 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>vpls>sap>egress>queue-override>queue
config>service>vpls>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 or default

high-prio-only

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

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>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>vpls>sap>egress>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.

Values 0 — 131072 or default

mbs

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

Context config>service>vpls>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's MBS parameters. The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packets 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 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 returns the MBS size assigned to the queue to the default value.

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

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled 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 — 100000000

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.

Values 0 — 100000000, **max**, **sum**

Default 0

scheduler-override

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

Context config>service>vpls>sap>egress
config>service>vpls>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>vpls>sap>egress>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>vpls>sap>egress>sched-override>scheduler

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 then 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 ate 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 queues 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**

scheduler-policy

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

Context config>service>vpls>sap>ingress
config>service>vpls>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.

match-qinq-dot1p

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

Context config>service>vpls>sap>ingress

Description This command specifies which Dot1Q tag position 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 16 defines the default behavior for Dot1P evaluation when the **match-qinq-dot1p** command is not executed.

Table 1: 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 17 defines the dot1p evaluation behavior when the top parameter is specified.

Table 2: 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 18 defines the dot1p evaluation behavior when the bottom parameter is specified.

Table 3: 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 3: 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 4: 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 5: 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

Table 5: QinQ Mark Top Only Explicit Marking Actions

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
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 PBits marked using dot1p-value, BottomQ PBits marked with zero
QinQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits PBits marked using dot1p-value, BottomQ PBits marked using preserved value

The QinQ and TopQ SAP PBit marking follows the default behavior devined in Table 19 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.

Service Billing Commands

authentication-policy

Syntax **authentication-policy** *name*
no authentication-policy

Context config>service>vpls>sap

Description This command defines which subscriber authentication policy must be applied when a DHCP message is received on the interface. The authentication policies must already be defined. The policy will only be applied when DHCP snooping is enabled on the SAP.

accounting-policy

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

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description This command creates the accounting policy context that can be applied to a SAP or SDP.

An accounting policy must be defined before it can be associated with a SAP or SDP. If the *policy-id* does not exist, an error message is generated.

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

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

Default Default accounting policy.

Parameters *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 — 99

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description This command enables accounting and statistical data collection for either the SAP or SDP, network port, or IP interface. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default **collect-stats**

VPLS SDP Commands

mesh-sdp

Syntax **mesh-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}]
no mesh-sdp *sdp-id*[:*vc-id*]

Context config>service>vpls

Description This command binds a VPLS service to an existing Service Distribution Point (SDP).

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

Note that this command creates a binding between a service and an SDP. The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate the SDP with a valid service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end 7750 SR devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

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

Special Cases **VPLS —** Several SDPs can be bound to a VPLS. Each SDP must be destined to a different 7750 SR router. If two *sdp-id* bindings terminate on the same 7750 SR, an error occurs and the second SDP is binding is rejected.

Parameters *sdp-id* — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit identifier. This value is used to validate the VC ID portion of each mesh SDP binding defined in the service. The default value of this object is equal to the service ID.

Values 1 — 4294967295

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

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

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

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

vlan — Defines the VC type as VLAN. The **ethernet**, **vlan,** and **vpls** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings.

vpls — Defines the VC type as VPLS. The **ethernet**, **vlan,** and **vpls** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings. Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke SDP binding.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*]
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>vpls

Description This command binds a service to an existing Service Distribution Point (SDP).

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

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a VPLS service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end 7750 SR devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

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

Special Cases **VPLS —** Several SDPs can be bound to a VPLS service. Each SDP must use unique *vc-ids*. An error message is generated if two SDP bindings with identical *vc-ids* terminate on the same router. Split horizon groups can only be created in the scope of a VPLS service.

Parameters *sdp-id* — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

vc-type — This command overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for

the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the bindings VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

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

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

Values ether, vlan

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

vlan — Defines the VC type as VLAN. The **ethernet**, **vlan,** and **vpls** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings. The VLAN VC-type requires at least one dot1Q tag within each encapsulated Ethernet packet transmitted to the far end.

split-horizon-group *group-name* **—** Specifies the name of the split horizon group to which the SDP belongs.

egress

Syntax **egress**

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command configures the egress SDP context.

ingress

Syntax **ingress**

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command configures the ingress SDP context.

mfib-allowed-mda-destinations

Syntax **mfib-allowed-mda-destinations**

Context config>service>vpls>mesh-sdp>egress
config>service>vpls>spoke-sdp>egress

Description This command enables the context to configure MFIB-allowed MDA destinations.

The allowed-mda-destinations node and the corresponding **mda** command are used on spoke and mesh SDP bindings to provide a list of MDA destinations in the chassis that are allowed as destinations for multicast streams represented by [*,g] and [s,g] multicast flooding records on the VPLS service. The MDA list only applies to IP multicast forwarding when IGMP snooping is enabled on the VPLS service. The MDA list has no effect on normal VPLS flooding such as broadcast, L2 multicast, unknown destinations or non-snooped IP multicast.

At the IGMP snooping level, a spoke or mesh SDP binding is included in the flooding domain for an IP multicast stream when it has either been defined as a multicast router port, received a IGMP query through the binding or has been associated with the multicast stream through an IGMP request by a host over the binding. Due to the dynamic nature of the way that a spoke or mesh SDP binding is associated with one or more egress network IP interfaces, the system treats the binding as appearing on all network ports. This causes all possible network destinations in the switch fabric to be included in the multicast streams flooding domain. The MDA destination list provides a simple mechanism that narrows the IP multicast switch fabric destinations for the spoke or mesh SDP binding.

If no MDAs are defined within the allowed-mda-destinations node, the system operates normally and will forward IP multicast flooded packets associated with the spoke or mesh SDP binding to all switch fabric taps containing network IP interfaces.

The MDA inclusion list should include all MDAs that the SDP binding may attempt to forward through. A simple way to ensure that an MDA that is not included in the list is not being used by the binding is to define the SDP the binding is associated with as MPLS and use an RSVP-TE LSP with a strict egress hop. The MDA associated with the IP interface defined as the strict egress hop should be present in the inclusion list.

If the inclusion list does not currently contain the MDA that the binding is forwarding through, the multicast packets will not reach the destination represented by the binding.

By default, the MDA inclusion list is empty.

If an MDA is removed from the list, the MDA is automatically removed from the flooding domain of any snooped IP multicast streams associated with a destination on the MDA unless the MDA was the last MDA on the inclusion list. Once the inclusion list is empty, all MDAs are eligible for snooped IP multicast flooding for streams associated with the SDP binding.

mda

Syntax [**no**] **mda** *mda-id*

Context config>service>vpls>mesh-sdp>egress>mfib-allowed-mda-destinations
config>service>vpls>spoke-sdp>egress>mfib-allowed-mda-destinations

Description This command specifies an MFIB-allowed MDA destination for an SDP binding configured in the system.

Parameters *mda-id* — Specifies an MFIB-allowed MDA destination.

Values

slot/mda	slot:	1 — 10
	mda:	1 — 2

vc-label

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

Context config>service>vpls>mesh-sdp>egress
config>service>vpls>spoke-sdp>egress

Description This command configures the egress VC label.

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

Values 16 — 1048575

vc-label

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

Context config>service>vpls>mesh-sdp>ingress
config>service>vpls>spoke-sdp>ingress

Description This command configures the ingress VC label.

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

Values 2048 — 18431

static-mac

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

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command creates a remote static MAC entry in the Virtual Private LAN Service (VPLS) forwarding database (FDB) associated with the Service Distribution Point (SDP).

In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Remote static MAC entries create a permanent MAC address to SDP association in the forwarding database for the VPLS instance so that MAC address will not be learned on the edge device.

Note that static MAC defnitions on one edge device are not propagated to other edge devices participating in the VPLS instance, that is, each edge device has an independent forwarding database for the VPLS.

Only one static MAC entry (local or remote) can be defined per MAC address per VPLS instance.

By default, no static MAC address entries are defined for the SDP.

The **no** form of this command deletes the static MAC entry with the specified MAC address associated with the SDP from the VPLS forwarding database.

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.

vlan-vc-tag

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

Context config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description This command specifies an explicit dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured dot1q tag can be overridden by a received TLV specifying the dot1q value expected by the far end. This signaled value must be stored as the remote signaled dot1q value for the binding. The provisioned local dot1q tag must be stored as the administrative dot1q value for the binding.

When the dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.

The **no** form of this command disables the command

Default **no vlan-vc-tag**

Parameters *0..4094* — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

root-guard

Syntax [**no**] **root-guard**

Context config>service>vpls>sap>stp

Description This command specifies whether this port is allowed to become an STP root port. It corresponds to the restrictedRole parameter in 802.1Q. If set, it can cause lack of spanning tree connectivity.

Default **no root-guard**

SAP Subscriber Management Commands

sub-sla-mgmt

Syntax [**no**] **sub-sla-mgmt**

Context config>service>vpls>sap

Description This command enables the context to configure subscriber management parameters for this SAP.

Default **no sub-sla-mgmt**

def-sla-profile

Syntax **def-sla-profile** *default-sla-profile-name*
no def-sla-profile

Context config>service>vpls>sap>sub-sla-mgmt

Description This command specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sla-profile** context.

An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts. SLA profiles are maintained in two locations, the subscriber identification policy and the subscriber profile templates. After a subscriber host is associated with an SLA profile name, either the subscriber identification policy used to identify the subscriber or the subscriber profile associated with the subscriber host must contain an SLA profile with that name. If both the subscriber identification policy and the subscriber profile contain the SLA profile name, the SLA profile in the subscriber profile is used.

The **no** form of the command removes the default SLA profile from the SAP configuration.

Default **no def-sla-profile**

Parameters *default-sla-profile-name* — Specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sla-profile** context.

def-sub-profile

Syntax **def-sub-profile** *default-subscriber-profile-name*

Context config>service>vpls>sap>sub-sla-mgmt

Description This command specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-profile** context.

A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for

subscriber using the subscriber profile. Subscriber profiles also allow for specific SLA profile definitions when the default definitions from the subscriber identification policy must be overridden.

The **no** form of the command removes the default SLA profile from the SAP configuration.

Parameters *default-sub-profile* — Specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-profile** context.

mac-da-hashing

Syntax [**no**] **mac-da-hashing**

Context config>service>vpls>sap>sub-sla-mgmt

Description This command specifies whether subscriber traffic egressing a LAG SAP has its egress LAG link selected by a function of the MAC destination address instead of the subscriber ID.

This command is only meaningful if subscriber management is enabled and can be configured for this VPLS service.

multi-sub-sap

Syntax **multi-sub-sap** [*subscriber-limit*]
no multi-sub-sap

Context config>service>vpls>sap>sub-sla-mgmt

Description This command configures the maximum number of subscribers for this SAP.

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

Parameters *number-of-sub* — Specifies the maximum number of subscribers for this SAP.

Values 2 — 8000

non-sub-traffic

Syntax **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*]
no non-sub-traffic

Context config>service>vpls>sap>sub-sla-mgmt>single-sub

Description This command configures non-subscriber traffic profiles.

The **no** form of the command removes the profiles and disables the feature.

Parameters **sub-profile** *sub-profile-name* **—** Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

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

subscriber *sub-ident-string* — Specifies an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

- For SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's sub-ident-string is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the service destinations.

  If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

  If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.

  If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

  ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

profiled-traffic-only

Syntax [**no**] **profiled-traffic-only**

Context config>service>vpls>sap>sub-sla-mgmt>single-sub

Description This command enables profiled traffic only for this SAP.

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

single-sub-parameters

Syntax **single-sub-parameters**

Context config>service>vpls>sap>sub-sla-mgmt

Description This command enables the context to configure single subscriber parameters for this SAP.

sub-ident-policy

Syntax **sub-ident-policy** *sub-ident-policy-name*

Context config>service>vpls>sap>sub-sla-mgmt

Description This command associates a subscriber identification policy to this SAP. The subscriber identification policy must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-ident-policy** context.

Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.

For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet string. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.

When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.

The **no** form of the command removes the default subscriber identifcation policy from the SAP configuration.

Default **no sub-ident-policy**

Parameters *sub-ident-policy-name* — Specifies a subscriber identification policy for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-ident-policy** context.

VPLS Multicast Commands

fast-leave

Syntax [**no**] **fast-leave**

Context config>service>vpls>sap>snooping
config>service>vpls>spoke-sdp>snooping
config>service>vpls>mesh-sdp>snooping

Description This command enables fast leave.

When IGMP fast leave processing is enabled, the SR-Series will immediately remove a SAP or SDP from the IP multicast group when it detects an IGMP 'leave' on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a 'leave' from the forwarding table without first sending out group-specific queries to the SAP or SDP, and thus speeds up the process of changing channels ('zapping').

Fast leave should only be enabled when there is a single receiver present on the SAP or SDP.

When fast leave is enabled, the configured last-member-query-interval value is ignored.

Default **no fast-leave**

from-vpls

Syntax **from-vpls** *vpls-id*
no from-vpls

Context config>service>vpls>sap>snooping>mvr

Description This command configures the VPLS from which multicast traffic is copied upon receipt of an IGMP join request.

IGMP snooping must be enabled on the MVR VPLS.

Default **no from-vpls**

Parameters *vpls-id* — Specifies the MVR VPLS from which multicast channels should be copied into this SAP.

group

Syntax [**no**] **group** *grp-address*

Context config>service>vpls>sap>snooping>static
config>service>vpls>spoke-sdp>snooping>static
config>service>vpls>mesh-sdp>snooping>static

Description This command adds a static multicast group either as a (*, g) or as one or more (s,g) records. When a static IGMP group is added, multicast data for that (*,g) or (s,g) is forwarded to the specific SAP or SDP without receiving any membership report from a host.

Default **none**

Parameters *grp-address* — Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group.

group-policy

Syntax **group-policy** *policy-name*
no group-policy

Context config>service>vpls>snooping>mvr

Description Identifies filter policy of multicast groups to be applied to this MVR VPLS. The sources of the multicast traffic must be a member of the MVR VPLS.

The **no** form of the command removes the MVR policy association from the VPLS.

Default No group policy is specified.

Parameters *policy-name* — The group 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. Routing policies are configured in the config>router>policy-options context. The router policy must be defined before it can be imported. For details on IGMP policies, please see section "Enabling IGMP group membership report filtering" in the *7750 SR OS Router Configuration Guide*.

igmp-snooping

Syntax **igmp-snooping**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description This command enables the Internet Group Management Protocol (IGMP) snooping context.

Default **none**

import

Syntax **import** *policy-name*
no import

Context config>service>vpls>sap>snooping
config>service>vpls>spoke-sdp>snooping
config>service>vpls>mesh-sdp>snooping

Description This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a single SAP at any time.

The **no** form of the command removes the policy association from the SAP or SDP.

Default **no import** — No import policy is specified.

Parameters *policy-name* — The import 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. Routing policies are configured in the config>router>policy-options context The router policy must be defined before it can be imported.

last-member-query-interval

Syntax **last-member-query-interval** *tenths-of-seconds*
no last-member-query-interval

Context config>service>vpls>sap>snooping
config>service>vpls>spoke-sdp>snooping
config>service>vpls>mesh-sdp>snooping

Description This command configures the maximum response time used in group-specific queries sent in response to 'leave'messages, and is also the amount of time between 2 consecutive group-specific queries. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.

The configured last-member-query-interval is ignored when fast-leave is enabled on the SAP or SDP.

Default **10**

Parameters *seconds* — Specifies the frequency, in tenths of seconds, at which query messages are sent.

Values 1 — 50

mcac

Syntax **mcac**

Context config>service>vpls>mesh-sdp>snooping
config>service>vpls>spoke-sdp>snooping
config>service>vpls>sap>snooping

Description This command configures multicast CAC policy and constraints for this interface.

Default **none**

policy

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

Context config>service>vpls>mesh-sdp>snooping>mcac
config>service>vpls>spoke-sdp>snooping>mcac
config>service>vpls>sap>snooping>mcac

Description This command configures the multicast CAC policy name.

Parameters *policy-name* — 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** *bandwidth* **mandatory-bw** *mandatory-bw*
no unconstrained-bw

Context config>service>vpls>mesh-sdp>snooping>mcac
config>service>vpls>spoke-sdp>snooping>mcac
config>service>vpls>sap>snooping>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 *bandwidth* — The bandwidth assigned for interface's MCAC policy traffic, in kilo-bits per second (kbps).

Values 0 — 2147483647

mandatory-bw *mandatory-bw* — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilo-bits per second (kbps).

If the *bandwidth* value is 0, no mandatory channels are allowed. If the value of *bandwidth* is '-1', then all mandatory and optional channels are allowed.

If the value of *mandatory-bw* is equal to the value of *bandwidth*, 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 *mandatory-bw* should always be less than or equal to that of *bandwidth*, An attempt to set the value of *mandatory-bw* greater than that of *bandwidth*, will result in inconsistent value error.

Values 0 — 2147483647

mc-constraints

Syntax **mc-constraints**

Context config>service>vpls>sap>snooping>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>vpls>sap>snooping>mcac>mc-constraints

Description This command configures levels and their associated bandwidth for multicast cac policy on this interface.

Parameters *level-id* **—** Specifies has an entry for each 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** *number-lag-port-down*
no number-down

Context config>service>vpls>sap>snooping>mcac>mc-constraints

Description This command configure the number of ports down along with level for multicast cac policy on this interface.

Default **not enabled**

max-num-groups

Syntax **max-num-groups** *count*
no max-num-groups

Context config>service>vpls>sap>snooping
config>service>vpls>spoke-sdp>snooping
config>service>vpls>mesh-sdp>snooping

Description This command defines the maximum number of multicast groups that can be joined on this SAP or SDP. If the SR-Series receives an IGMP join message that would exceed the configured number of groups, the request is ignored.

Default **no max-num-groups**

Parameters *count* — Specifies the maximum number of groups that can be joined on this SAP or SDP.

Values 1 — 1000

mrouter-port

Syntax [**no**] **mrouter-port**

Context config>service>vpls>sap>snooping
config>service>vpls>spoke-sdp>snooping
config>service>vpls>mesh-sdp>snooping

Description This command specifies whether a multicast router is attached behind this SAP or SDP.

Configuring a SAP or SDP as an mrouter-port will have a double effect. Firstly, all multicast traffic received on another SAP or SDP will be copied to this SAP or SDP. Secondly, IGMP reports generated by the system as a result of someone joining or leaving a multicast group, will be sent to this SAP or SDP.

If two multicast routers exist in the network, one of them will become the active querier. While the other multicast router (non-querier) stops sending IGMP queries, it should still receive reports to keep its multicast trees up to date. To support this, the mrouter-port should be enabled on all SAPs or SDPs connecting to a multicast router.

Note that the IGMP version to be used for the reports (v1, v2 or v3) can only be determined after an initial query has been received. Until such time no reports are sent on the SAP or spoke SDP, even if mrouter-port is enabled.

This parameter can only be enabled on a SAP or spoke SDP, not on a mesh SDP.

If the **send-queries** command is enabled on this SAP or spoke SDP, the mrouter-port parameter can not be set.

Default **no mrouter-port**

mvr

Syntax **mvr**

Context config>service>vpls>snooping
config>service>vpls>sap>snooping

Description This command enables the context to configure Multicast VPLS Registration (MVR) parameters.

query-interval

Syntax **query-interval** *seconds*
no query-interval

Context config>service>vpls>snooping
config>service>vpls>sap>snooping
config>service>vpls>spoke-sdp>snooping
config>service>vpls>mesh-sdp>snooping

Description If the **send-queries** command is enabled, this parameter specifies the interval between two consecutive general queries sent by the system on this SAP or SDP.

The configured query-interval must be greater than the configured query-response-interval.

If send-queries is not enabled on this SAP or SDP, the configured query-interval value is ignored.

Default **125**

Parameters *seconds* — The time interval, in seconds, that the router transmits general host-query messages.

Values 2 — 1024

query-response-interval

Syntax **query-response-interval** *seconds*

Context config>service>vpls>sap>snooping
config>service>vpls>spoke-sdp>snooping
config>service>vpls>mesh-sdp>snooping

Description If the **send-queries** command is enabled, this parameter specifies the maximum response time advertised in IGMPv2/v3 queries.

The configured query-response-interval must be smaller than the configured query-interval.

If send-queries is not enabled on this SAP or SDP, the configured query-response-interval value is ignored.

Default **10**

Parameters *seconds* — Specifies the length of time to wait to receive a response to the host-query message from the host.

Values 1 — 1023

report-src-ip

Syntax **report-src-ip**-*address*
no report-src-ip

Context config>service>vpls>igmp-snooping

Description This parameter specifies the source IP address used when generating IGMP reports. According the IGMPv3 standard, a zero source address is allowed in sending IGMP reports. However, for interoperability with some multicast routers, the source IP address of IGMP group reports can be configured using this command.

Default **0.0.0.0**

Parameters *ip-address* — The source IP source address in transmitted IGMP reports.

robust-count

Syntax **robust-count** *robust-count*
no robust-count

Context config>service>vpls>snooping
config>service>vpls>sap>snooping
config>service>vpls>spoke-sdp>snooping
config>service>vpls>mesh-sdp>snooping

Description If the **send-queries** command is enabled, this parameter allows tuning for the expected packet loss on a SAP or SDP. The robust-count variable allows tuning for the expected packet loss on a subnet and is comparable to a retry count. If this SAP or SDP is expected to be 'lossy', this parameter may be increased. IGMP snooping on this SAP or SDP is robust to (robust-count-1) packet losses.

If send-queries is not enabled, this parameter will be ignored.

Default **2**

Parameters *robust-count* — Specifies the robust count for the SAP or SDP.

Values 2 — 7

send-queries

Syntax [**no**] **send-queries**

Context config>service>vpls>sap>snooping
config>service>vpls>spoke-sdp>snooping
config>service>vpls>mesh-sdp>snooping

Description This command specifies whether to send IGMP general query messages on the SAP or SDP.

When send-queries is configured, all type of queries generate ourselves are of the configured version. If a report of a version higher than the configured version is received, the report will get dropped and a new wrong version counter will get incremented.

If send-queries is not configured, the version command has no effect. The version used on that SAP/SDP will be the version of the querier. This implies that, for example, when we have a v2 querier, we will never send out a v3 group or group-source specific query when a host wants to leave a certain group.

Default **no send-queries**

source

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

Context config>service>vpls>sap>snooping>static>group
config>service>vpls>spoke-sdp>snooping>static>group
config>service>vpls>mesh-sdp>snooping>static>group

Description This command adds a static (s,g) entry to allow multicast traffic for the corresponding multicast group from that specific source. For the same multicast group, more than one source can be specified.

Static (s,g) entries can not be entered when a starg is already created.

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 [**no**] **starg**

Context config>service>vpls>sap>snooping>static>group
config>service>vpls>spoke-sdp>snooping>static>group
config>service>vpls>mesh-sdp>snooping>static>group

Description This command adds a static (*,g) entry to allow multicast traffic for the corresponding multicast group from any source. This command can only be enabled if 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 **no starg**

static

Syntax **static**

Context config>service>vpls>sap>snooping
config>service>vpls>spoke-sdp>snooping
config>service>vpls>mesh-sdp>snooping

Description This command enables access to the context to configure static group addresses. Static group addresses can be configured on a SAP or SDP. When present either as a (*, g) or a (s,g) entry, multicast packets matching the configuration will be forwarded even if no join message was registered for the specific group.

Default **none**

version

Syntax **version** *version*
no version

Context config>service>vpls>sap>snooping
config>service>vpls>mesh-sdp>snooping
config>service>vpls>spoke-sdp>snooping

Description This command specifies the version of IGMP which is running on this SDP or SAP. This object can be used to configure a router capable of running either value. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN.

When the **send-query** command is configured, all type of queries generate ourselves are of the configured **version**. If a report of a version higher than the configured version is received, the report gets dropped and a new "wrong version" counter is incremented.

If the **send-query** command is not configured, the **version** command has no effect. The version used on that SAP or SDP will be the version of the querier. This implies that, for example, when there is a v2 querier, a v3 group or group-source specific query when a host wants to leave a certain group will never be sent.

Parameters *version* — Specify the IGMP version.

Values 1, 2, 3

to-sap

Syntax **to-sap** *sap-id*
no to-sap

Context config>service>vpls>sap>snooping>mvr

Description In some situations, the multicast traffic should not be copied from the MVR VPLS to the SAP on which the IGMP message was received (standard MVR behaviour) but to another SAP.

This command configures the SAP to which the multicast data needs to be copied.

Default **no to-sap**

Parameters *sap-id* — Specifies the SAP to which multicast channels should be copied.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type*-*slot/mda*.*bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type*-*bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094

lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

VPLS DHCP and Anti-Spoofing Commands

anti-spoof

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

Context config>service>vpls>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**, **mac**, **ip-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 ip** command will fail.

mac — Configures SAP anti-spoof filtering to use only the source MAC address in its lookup. If a static host exists on the SAP without a specified MAC address, the **anti-spoof mac** command will fail.

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 ip-mac** command will fail.

arp-reply-agent

Syntax **arp-reply-agent** [**sub-ident**]
no arp-reply-agent

Context config>service>vpls>sap

Description This command enables a special ARP response mechanism in the system for ARP requests destined to static or dynamic hosts associated with the SAP. The system responds to each ARP request using the hosts MAC address as the both the source MAC address in the Ethernet header and the target hardware address in the ARP header.

ARP replies and requests received on a SAP with **arp-reply-agent** enabled will be evaluated by the system against the anti-spoof filter entries associated with the ingress SAP (if the SAP has anti-spoof filtering enabled). ARPs from unknown hosts on the SAP will be discarded when anti-spoof filtering is enabled.

The ARP reply agent only responds if the ARP request enters an interface (SAP, spoke-SDP or mesh-SDP) associated with the VPLS instance of the SAP.

A received ARP request that is not in the ARP reply agent table is flooded to all forwarding interfaces of the VPLS capable of broadcast except the ingress interface while honoring split-horizon constraints.

Static hosts can be defined on the SAP using the **host** command. Dynamic hosts are enabled on the system by enabling the **lease-populate** command in the SAP's **dhcp** context. In the event that both a static host and a dynamic host share the same IP and MAC address, the VPLS ARP reply agent will retain the host information until both the static and dynamic information are removed. In the event that both a static and dynamic host share the same IP address, but different MAC addresses, the VPLS ARP reply agent is populated with the static host information.

The **arp-reply-agent** command will fail if an existing static host on the SAP does not have both MAC and IP addresses specified. Once the ARP reply agent is enabled, creating a static host on the SAP without both an IP address and MAC address will fail.

The ARP-reply-agent may only be enabled on SAPs supporting Ethernet encapsulation.

The **no** form of the command disables ARP-reply-agent functions for static and dynamic hosts on the SAP.

Default **not enabled**

Parameters **sub-ident** — Configures the arp-reply-agent to discard ARP requests received on the SAP that are targeted for a known host on the same SAP with the same subscriber identification.

Hosts are identified by their subscriber information. For DHCP subscriber hosts, the subscriber hosts, the subscriber information is configured using the optional subscriber parameter string.

When arp-reply-agent is enabled with **sub-ident**:

- If the subscriber information for the destination host exactly matches the subscriber information for the originating host and the destination host is known on the same SAP as the source, the ARP request is silently discarded.
- If the subscriber information for the destination host or originating host is unknown or undefined, the source and destination hosts are not considered to be the same subscriber. The ARP request is forwarded outside the SAPs Split Horizon Group.
- When **sub-ident** is not configured, the arp-reply-agent does not attempt to identify the subscriber information for the destination or originating host and will not discard an ARP request based on subscriber information.

host

Syntax **host** {[**ip** *ip-address*] [**mac** *ieee-address*]}[**subscriber** *sub-ident-string*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**ancp-string** *ancp-string*]
no host {[**ip** *ip-address*] [**mac** *ieee-address*]}
no host all

Context config>service>vpls>sap

Description This command creates a static subscriber host for the SAP. Static subscriber hosts may be used by the system for various purposes. Applications within the system that make use of static host entries include anti-spoof, ARP reply agent and source MAC population into the VPLS forwarding database.

Multiple static hosts may be defined on the SAP. Each host is identified by either a source IP address, a source MAC address or both a source IP and source MAC address. Every static host definition must have at least one address defined, IP or MAC.

Static hosts can exist on the SAP even with anti-spoof and ARP reply agent features disabled. When enabled, each feature has different requirements for static hosts.

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

Default **none**

Parameters **ip** *ip-address* — Specify this optional parameter when defining a static host. The IP address must be specified for **anti-spoof ip** and **anti-spoof ip-mac**. Only one static host may be configured on the SAP with a given IP address.

A static host with an ip-only definition is supported and all respective limitations (such as, enabling ip-mac and antispoof at the same time) are removed.

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.

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

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

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

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

host-connectivity-verify

Syntax **host-connectivity-verify source-ip** *ip-address* [**source-mac** *ieee-address*] [**interval** *interval*] [**action** {**remove** | **alarm**}]

Context config>service>vpls
config>service>vpls>sap

Description This command enables enables subscriber host connectivity verification on a given SAP within a VPLS service. This tool will periodically scan all known hosts (from dhcp-state) and perform a UC ARP request. The subscriber host connectivity verification will maintain state (connected vs. not-connected) for all hosts.

Default **no host-connectivity-verify**

Parameters **source-ip** *ip-address* **—** Specifies the source IP address to be used for generation of subscriber host connectivity verification packets.

source-mac *ieee-address* **—** Specifies the source MAC address to be used for generation of subscriber host connectivity verification packets.

interval ***interval*** **—** The interval, in minutes, which specifies the time interval which all known sources should be verified. The actual rate is then dependent on number of known hosts and interval.

Values 1 — 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, etc.). DHCP release will be signaled to corresponding DHCP server. Static host will be never removed. The **alarm** keyword raises an alarm indicating that the host is disconnected.

Egress Multicast Group Commands

egress-multicast-group

Syntax **egress-multicast-group** *egress-multicast-group-name*
no egress-multicast-group *group-name*

Context config>service

Description This command creates an egress multicast group (EMG) context. An EMG is created as an object used to group VPLS SAPs that are allowed to participate in efficient multicast replication (EMR). EMR is a method to increase the performance of egress multipoint forwarding by sacrificing some destination-based features. Eliminating the requirement to perform unique features for each destination allows the egress forwarding plane to chain together multiple destinations into a batch replication process. In order to perform this batch replication function, similar characteristics are required on each SAP within the EMG.

Only SAPs defined on Ethernet access ports are allowed into an egress-multicast-group.

In order to understand the purpose of an egress-multicast-group, an understanding of the system's use of flooding lists is required. A flooding list is maintained at the egress forwarding plane to define a set of destinations to which a packet must be replicated. Multipoint services make use of flooding lists to enable forwarding a single packet to many destinations. Examples of multipoint services that use flooding lists are VPLS, IGMP snooping and IP multicast routing. Currently, the egress forwarding plane will only use efficient multicast replication for VPLS and IGMP snooping flooding lists.

In VPLS services, a unique flooding list is created for each VPLS context. The flooding list is used when a packet has a broadcast, multicast or unknown destination MAC address. From a system perspective, proper VPLS handling requires that a broadcast, multicast or unknown destined packet be sent to all destinations that are in the forwarding state. The ingress forwarding plane ensures the packet gets to all egress forwarding planes that include a destination in the VPLS context. It is the egress forwarding plane's job to replicate the packet to the subset of the destinations that are reached through its interfaces and each of these destinations are included in the VPLS context's flooding list.

For IGMP snooping, a unique flooding list is created for each IP multicast (s,g) record. This (s,g) record is associated with an ingress VPLS context and may be associated with VPLS destinations in the source VPLS instance or other VPLS instances (in the case of MVR). Again, the ingress forwarding plane ensures that an ingress IP multicast packet matching the (s,g) record gets to all egress forwarding planes that have a VPLS destination associated with the (s,g) record. The egress forwarding plane uses the flooding list owned by the (s,g) record to replicate the packet to all VPLS destinations in the flooding list. The IGMP Snooping function identifies which VPLS destinations should be associated with the (s,g) record.

With normal multicast replication, the egress forwarding plane examines which features are enabled for each destination. This includes ACL filtering, mirroring, encapsulation and queuing. The resources used to perform this per destination multicast processing are very expensive to the egress forwarding plane when high replication bandwidth is required. If destinations with similar egress functions can be grouped together, the egress forwarding plane can process them in a more efficient manner and maximize replication bandwidth.

The egress-multicast-group object is designed to allow the identification of SAPs with similar egress characteristics. When a SAP is successfully provisioned into an egress-multicast-group, the system is ensured that it may be batched together with other SAPs in the same group at the egress forwarding

plane for efficient multicast replication. A SAP that does not meet the common requirements is not allowed into the egress-multicast-group.

At the forwarding plane level, a VPLS flooding list is categorized into chainable and non-chainable destinations. Currently, the only chainable destinations are SAPs within an egress-multicast-group. The chainable destinations are further separated by egress-multicast-group association. Chains are then created following the rules below:

- A replication batch chain may only contain SAPs from the same egress-multicast-group
- A replication batch chain length may not exceed the dest-chain-limit of the egress-multicast-group to which the SAPs are members

Further subcategories are created for an IGMP (s,g) flooding list. A Layer 2 (s,g) record is created in a specific VPLS instance (the instance the (s,g) flow ingresses). SAPs within that VPLS context that join the (s,g) record are considered native SAPs within the flooding list. SAPs that join the (s,g) flooding list through the multicast VPLS registration process (MVR) from another VPLS context using the **from-vpls** command are considered alien SAPs. The distinction between native and alien in the list is maintained to allow the forwarding plane to enforce or suspend split-horizon-group (SHG) squelching. When the source of the (s,g) matching packet is in the same SHG as a native SAP, the packet must not be replicated to that SAP. For a SAP in another VPLS context, the source SHG of the packet has no meaning and the forwarding plane must disregard SHG matching between the native source of the packet and the alien destination. Because the SHG squelch decision is done for the whole chain based on the first SAP in the chain, all SAPs in the chain must be all native or all alien SAPs. Chains for IGMP (s,g) flooding lists are created using the following rules:

1. A replication batch chain may only contain SAPs from the same egress-multicast-group.
2. A replication batch chain may only contain all alien or all native SAPs.
3. A replication batch chain length may not exceed the dest-chain-limit of the egress-multicast-group to which the SAPs are members

When a packet associated with a flooding list is received by the egress forwarding plane, it processes the packet by evaluating each destination on the list sequentially in a replication context. If the current entry being processed in the list is a non-chained destination, the forwarding plane processes the packet for that destination and then moves on to process other packets currently in the forwarding plane before returning to process the next destination in the list. If the current entry being processed is a chained destination, the forwarding plane remains in the replication context until it has forwarded to each entry in that chain. Once the replication context finishes with the last entry in the chain, it moves on to process other packets waiting for egress processing before returning to the replication context. Processing continues in this manner until the packet has been forwarded to all destinations in the list.

Batch chain processing of a chain of SAPs improves replication efficiency by bypassing the functions that perform egress mirroring decisions on SAPs within the chain and making a single ACL filtering decision for the whole chain. Each destination in the chain may have a unique egress QoS policy and per destination queuing is still performed for each destination in the chain. Also, while each SAP in the chain must be on access ports with the same encap-type, if the encap-type is dot1q, each SAP may have a unique dot1q tag.

One caveat to each SAP having a unique egress QoS policy in the chain is that only the Dot1P marking decisions for the first SAP in the list is enforced. If the first SAP's QoS policy forwarding class action states that the packet should not be remarked, none of the replicated packets in the chain will have the dot1P bits remarked. If the first SAP's QoS policy forwarding class action states that the packet should be remarked with a specific dot1P value, all the replicated packets for the remaining SAPs in the chain will have the same dot1P marking.

While the system supports 32 egress multicast groups, a single group would usually suffice. An instance where multiple groups would be needed is when all the SAPs requiring efficient multicast

replication cannot share the same common requirements. In this case, an egress multicast group would be created for each set of common requirements. An egress multicast group may contain SAPs from many different VPLS instances. It should be understood that an egress multicast group is not equivalent to an egress forwarding plane flooding list. An egress multicast group only identifies which SAPs may participate in efficient multicast replication. As stated above, entries in a flooding list are populated due to VPLS destination creation or IGMP snooping events.

The **no** form of the command removes a specific egress multicast group. Deleting an egress multicast group will only succeed when the group has no SAP members. To remove SAP members, use the **no multicast-group** *group-name* command under each SAP's egress context.

Note: Efficient multicast replication will only be performed on IOMs that support chassis mode b If an IOM does not support mode b operation, egress-multicast-group membership is ignored on that IOM's egress forwarding planes. The chassis need not be placed into mode b for efficient multicast replication to be performed on the capable IOMs.

Parameters *group-name* — Multiple egress multicast groups may be created on the system. Each must have a unique name. The egress-multicast-group-name is an ASCII string up to 16 characters in length and follows all the naming rules as other named policies in the system. The group's name is used throughout the system to uniquely identify the Egress Multicast Group and is used to provision a SAP into the group.

Default None, each egress multicast group must be explicitly configured.

Values Up to 32 egress multicast groups may be created on the system.

description

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

Context config>service>egress-multicast-group

Description This command defines an ASCII string associated with egress-multicast-group-name.

The **no** form of the command removes an existing description string from egress-multicast-group.

Default **none**

Parameters *description-string* — The description command accepts a description-string parameter. The description-string parameter is an ASCII string of up to 80 characters in length. Only printable 127 bit ASCII characters are allowed. If the string contains spaces, the string must be specified with beginning and ending quotes.

Values An ASCII string up to 80 characters in length.

dest-chain-limit

Syntax **dest-chain-limit** *destinations per pass*
no dest-chain-limit

Context config>service>egress-multicast-group

Description This command defines the maximum length of an egress forwarding plane efficient multicast replication chain for an egress-multicast-group. Varying the maximum length of chains created for an egress multicast group has the effect of efficient multicast batched chain replication on other packets flowing through the egress forwarding plane. While replicating for the SAPs within a replication chain, other packets are waiting for the forwarding plane to finish. As the chain length increases, forwarding latency for the other waiting packets may increase. When the chain length decreases, a loss of efficiency in the replication process will be observed.

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

Default **16**

Parameters *destinations per pass* — This parameter must be specified when executing the **dest-chain-limit** command. When executed, the command will use the number-of-destinations parameter to reorganize all efficient multicast SAP chains that contain members from the egress-multicast-group.

The *destinations per pass* parameter can be modified at any time. Be aware that when changing the maximum chain length, the system will rebuild the chains according to the new limit. When this happens, it is possible that packets will not be replicated to a destination while it is being reorganized in the flooding list's chains. Only the chains associated with the egress-multicast-group context the command is executed in will be affected by changing the parameter.

It is expected that the optimal replication chain length will be between 10 and 16. Since so many variables affect efficient multicast (i.e. ingress packet rate, number of chains, size of replicated packets), only proper testing in the environment that replication will be performed will identify the best dest-chain-limit value for each Egress Multicast Group.

Setting the *destinations per pass* parameter to a value of 0 has the effect of removing from all egress forwarding planes all chains with members from the egress-multicast-group. Replication to each destination SAP from the group is performed using the normal method (non-efficient replication). The value 0 is not considered a normal value for dest-chain-limit and is provided for debugging purposes only. Setting the value to 0 is persistent between reboots of the system.

Setting the *destinations per pass* parameter to a value of 1 has the effect of placing each egress-multicast-group member SAP into a chain with a single SAP. The value 1 is not considered a normal value for the **dest-chain-limit** and is provided for debugging purposes only. Setting the value to 1 is persistent between reboots of the system.

Values 1 — 30

sap-common-requirements

Syntax **sap-common-requirements**

Context config>service>egress-multicast-group

Description This command configures the common SAP parameter requirements. The SAP common requirements are used to evaluate each SAP for group membership. If a SAP does not meet the specified requirements, the SAP is not allowed into the egress-multicast-group. Once a SAP is a member of the group, attempting to change the parameters on the SAP will fail.

egress-filter

Syntax **egress-filter** [**ip** *ip-filter-id*]
egress-filter [**ipv6** *ipv6-filter-id*]
egress-filter [**mac** *mac-filter-id*]
no egress-filter [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*] [**mac** *mac-filter-id*]

Context config>service>egress-multicast-group>sap-common-requirements

Description This command identifies the type of filter and actual filter ID that must be provisioned on the SAP prior to the SAP being made a member of the egress-multicast-group. If the SAP does not have the specified filter applied, the SAP cannot be provisioned into the group. It is important that the egress filter applied to each SAP within the egress-multicast-group be the same since the batch replication process on an efficient multicast replication chain will apply the first SAP's ACL decision to all other SAPs on the chain.

Once the SAP is made a member of the egress-multicast-group, the SAP's egress filter cannot be changed on the SAP.

Changing the **egress-filter** parameters within the **sap-common-requirements** node automatically changes the egress filter applied to each member SAP. If the filter cannot be changed on the SAP due to resource constraints, the modification will fail.

The specified egress-filter does not contain an entry that is defined as an egress mirror-source. Once the filter is associated with the egress-multicast-group, attempting to define one of its entries as an egress mirror source will fail.

The **no** form of the command removes the egress-filter removes the egress filter from each member SAP. The **no egress-filter** command specifies that an egress filter (IP or MAC)(IP, IPv6 or MAC) is not applied to a new member SAP within the egress-multicast-group.

Default **no filter**. The egress filter ID must be defined with the associated **ip** or **mac** keyword. If an egress-filter is not specified or the no egress-filter command is executed in the sap-common-requirements node, a new member SAP does not have an egress IP or MAC filter defined.

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

Values 1 — 65535

ipv6 *ipv6-filter-id* **—** Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 — 65535

mac *mac-filter-id* **—** Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1 — 65535

encap-type

Syntax **encap-type** {**dot1q** | **null**}
no encap-type

Context config>service>egress-multicast-group>sap-common-requirements

Description This command specifies the encapsulation type that must exist on the SAP's access port to allow the SAP membership within the egress-multicast-group. The **config>port>ethernet>access>encap-type** command is used to define the encapsulation type for the Ethernet port. The allowed encapsulation type values are dot1q and null. If the SAP does not exist on a port with the specified encap-type, it will not be allowed into the egress-multicast-group.

If at least one SAP is currently a member of the efficient-multicast-group, the **encap-type** cannot be changed within the sap-common-requirements node. If the efficient-multicast-group does not contain any member SAPs, the **encap-type** may be changed at any time.

There is no interaction between an efficient-multicast-group and the corresponding access ports associated with its members since all SAPs must be deleted from a port before its encap-type can be changed. When the SAPs are deleted from the port, they are also automatically deleted from the efficient-multicast-group.

The **no** form of the command returns the egress-multicast-group required encapsulation type for SAPs to dot1q. If the current encap-type is set to null, the command cannot be executed when SAPs exist within the egress-multicast-group.

Default **dot1q** — For an egress-multicast-group.
null — If member SAPs are on a null encapsulated access port.

Parameters **null —** The **null** keyword is mutually exclusive with the **dot1q** keyword. When the encap-type within the sap-common-requirements is specified to be null, the encapsulation type for the access ports associated with all SAPs within the egress-multicast-group must be set to null.

dot1q — The **dot1q** keyword is mutually exclusive with the **null** keyword. When the encap-type within the sap-common-requirements is specified to be dot1q, the encapsulation type for the access ports associated with all SAPs within the egress-multicast-group must be set to dot1q.

dot1q-etype

Syntax **dot1q-etype** [0x0600..0xffff]
no dot1q-etype

Context config>service>egress-multicast-group>sap-common-requirements

Description This command specifies the dot1q EtherType that must exist on the SAP's access port to allow the SAP membership within the egress-multicast-group. The **config>port>ethernet>access>dot1q-etype** command is used to define the EtherType used when encapsulating a packet with a dot1q tag on the Ethernet port. Any valid EtherType is allowed on the port.

If the current encap-type for the egress-multicast-group is set to null, the dot1q-etype EtherType is ignored when evaluating SAP membership in the group. If the encap-type is set to dot1q (the default), a member SAP's access port must be configured with the same dot1q-etype EtherType as the egress-multicast-group.

If at least one SAP is currently a member of the efficient-multicast-group, the dot1q-etype value cannot be changed within the sap-common-requirements node. If the efficient-multicast-group does not contain any member SAPs, the dot1q-etype value may be changed at any time.

If an access port currently has SAPs associated with it that are defined within an egress-multicast-group and the port is currently set to encap-type dot1q, the dot1q-etype value defined on the port cannot be changed.

The **no** form of the command returns the egress-multicast-group dot1q EtherType to the default value of 0x8100. If the current encap-type is set to a value other then 0x8100, the command cannot be executed when SAPs exist within the egress-multicast-group.

Default The default dot1q-etype is 0x8100 for an egress-multicast-group.

Parameters *ethertype* — Defines the dot1q EtherType that must be associated with a SAP's access port when the encap-type is set to dot1q. Any valid EtherType may be specified.

Values [0x0600 — 0xffff]: [1536 — 65535] in decimal or hex

Default 0x8100

Provider Edge Discovery Policy Commands

pe-discovery-policy

Syntax [**no**] **pe-discovery-policy** *name*

Context config>service

Description This command configures a provider edge discovery policy and parameters.

The **no** form of the command removes the policy from the configureation.

Parameters *name* — Specifies theRADIUS PE discovery policy name, up to 32 characters in length.

password

Syntax **password** *password*
no password

Context config>service>pe-discovery-policy

Description This command configures the PE discovery password that is used when contacting the RADIUS server for VPLS auto-discovery.

The **no** form of the command removes the password from the configuration.

Default **no password**

Parameters *password* — Specifies the password, up to 32 characters in length, used when contacting the RADIUS server for VPLS auto-discovery.

polling-interval

Syntax **polling-interval** *minutes*
no polling-interval

Context config>service>pe-discovery-policy

Description This command configures the PE discovery polling interval.

Default **5**

Parameters *minutes* — Specifies the polling interval, in minutes, for RADIUS PE discovery.

Values 1 — 30

server

Syntax **server** *server-index* **address** *ip-address* **secret** *key* [**hash** | **hash2**] [**port** *port-num*]
no server *server-index*

Context config>service>pe-discovery-policy

Description This command adds a RADIUS server.

Parameters *server-index* — The index for the RADIUS server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index.

Values 1 — 5

address *ip-address* **—** The IP address of the RADIUS server. Two RADIUS servers cannot have the same IP address. An error message is generated if the server address is a duplicate.

secret *key* **—** The secret key to access the RADIUS server. This secret key must match the password on the RADIUS server.

Values Up to 20 characters in length.

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.

port — Specifies the UDP port number on which to contact the RADIUS server for authentication.

Values 1 — 65535

timeout

Syntax **timeout** *seconds*
no timeout

Context config>service>pe-discovery-policy

Description This command specifies the number of seconds to wait before timing out a RADIUS server.

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

Default **3 seconds**

Parameters *seconds* — The number of seconds to wait for a response from a RADIUS server, expressed as a decimal integer.

Values 1 — 90

radius-discovery

Syntax [**no**] **radius-discovery**

Context config>service>vpls

Description This command enables the RADIUS provider edge discovery for this VPLS service.

Default **none**

pe-discovery-policy

Syntax **pe-discovery-policy** *name*
no pe-discovery-policy

Context config>service>vpls>radius-discovery

Description This command specifies the existing RADIUS PE discovery policy name. The policy must have been configured in the **config**>**service** context.

Parameters *name* — Specifies the RADIUS PE discovery policy name, up to 32 characters in length.

user-name-format

Syntax **user-name-format** {**vpn-id** *vpn-id* | **router-distinguisher** *rd*}
no pe-discovery-policy

Context config>service>vpls>radius-discovery

Description This command specifies whether the RADIUS user name is a VPN ID or router-distinguisher.

Parameters **vpn-id** *vpn-id* — Indicates the VPN ID of the associated VPLS service.

router-distinguisher *rd* — Sets the identifier attached to routes that distinguishes the VPN it belongs.

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.

Table 6: 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#
```

fdb-info

Syntax **fdb-info**

Context show>service

Description Displays global FDB usage information.

Output **Show FDB-Info Command Output —** The following table describes show FDB-Info command output.

Label	Description
Service ID	The ID that identifies a service.
Mac Move	Indicates the administrative state of the MAC movement feature associated with the service.
Mac Move Rate	The maximum rate at which MAC's can be re-learned in this TLS service, before the SAP where the moving MAC was last seen is automatically disabled in order to protect the system against undetected loops or duplicate MAC's. The rate is computed as the maximum number of re-learns allowed in a 5 second interval: e.g. the default rate of 10 re-learns per second corresponds to 50 re-learns in a 5 second period.

Mac Move Timeout	Indicates the time in seconds to wait before a SAP that has been disabled after exceeding the maximum re-learn rate is re-enabled. A value of zero indicates that the SAP will not be automatically re-enabled after being disabled. If after the SAP is re-enabled it is disabled again, the effective retry timeout is doubled in order to avoid thrashing.
Table Size	The maximum number of learned and static entries allowed in the FDB of this service.
Total Count	The current number of entries (both learned and static) in the FDB of this service.
Learned Count	The current number of learned entries in the FDB of this service.
Static Count	The current number of static entries in the FDB of this service.
Remote Age	The number of seconds used to age out FDB entries learned on an SDP. These entries correspond to MAC addresses learned on remote SAPs.
Local Age	The number of seconds used to age out FDB entries learned on local SAPs.
High WaterMark	The utilization of the FDB table of this service at which a 'table full' alarm is raised by the agent.
Low WaterMark	The utilization of the FDB table of this service at which a 'table full' alarm is cleared by the agent.
Mac Learning	Specifies whether the MAC learning process is enabled in this service.
Discard Unknown	Specifies whether frames received with an unknown destination MAC are discarded in this service.
MAC Aging	Specifies whether the MAC aging process is enabled in this service.
Relearn Only	When enabled, indicates that either the FDB table of this service is full or that the maximum system-wide number of MAC's supported by the agent has been reached, and thus MAC learning is temporary disabled, and only MAC re-learns can take place.
Total Service FDBs	The current number of service FDBs configured on this node.
Total FDB Configured Size	The sum of configured FDBs.
Total FDB Entries In Use	The total number of entries (both learned and static) in use.

Sample Output

```
*A:ALA-12# show service fdb-info
===============================================================================
Forwarding Database (FDB) Information
===============================================================================
Service Id        : 700                Mac Move          : Disabled
Mac Move Rate     : 10                 Mac Move Timeout  : 10
```

```
Table Size        : 250                Total Count       : 0
Learned Count     : 0                  Static Count      : 0
Remote Age        : 900                Local Age         : 300
High WaterMark    : 95%                Low Watermark     : 90%
Mac Aging         : Enabl              Relearn Only      : False

Service Id        : 725                Mac Move          : Disabled
Mac Move Rate     : 10                 Mac Move Timeout  : 10
Table Size        : 250                Total Count       : 0
Learned Count     : 0                  Static Count      : 0
Remote Age        : 900                Local Age         : 300
High WaterMark    : 95%                Low Watermark     : 90%
Mac Learning      : Enabl              Discard Unknown   : Dsabl
Mac Aging         : Enabl              Relearn Only      : False

Service Id        : 740                Mac Move          : Disabled
Mac Move Rate     : 10                 Mac Move Timeout  : 10
Table Size        : 250                Total Count       : 0
Learned Count     : 0                  Static Count      : 0
Remote Age        : 900                Local Age         : 300
High WaterMark    : 95%                Low Watermark     : 90%
Mac Learning      : Enabl              Discard Unknown   : Dsabl
Mac Aging         : Enabl              Relearn Only      : False
...
-------------------------------------------------------------------------------
Total Service FDBs : 7
Total FDB Configured Size : 1750
Total FDB Entries In Use  : 0
-------------------------------------------------------------------------------
===============================================================================
A:*A:ALA-48#
```

fdb-mac

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

Context show>service

Description Displays the FDB entry for a given MAC address.

Parameters *ieee-address* — The 48-bit MAC address for which to display the FDB entry in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

expiry — Shows the time until the MAC is aged out.

Output **Show FDB-MAC Command Output —** the following table describes the show FDB MAC command 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.

Sample Output

```
*A:ALA-12# show service fdb-mac 00:99:00:00:00:00
===============================================================================
Services Using Forwarding Database Mac 00:99:00:00:00:00
===============================================================================
ServId  MAC                    Source-Identifier       Type/Age Last Change
-------------------------------------------------------------------------------
1       00:99:00:00:00:00             sap:1/2/7:0               Static
===============================================================================
*A:ALA-12#
```

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** *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 2049 — 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.

Label	Description
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** [**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 *auth-plcy-name*

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 **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 *auth-plcy-name* **—** The session authentication policy for which to display matching SAPs.

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]

	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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 service identifier.
SapMTU	The SAP MTU value.
I.QoS	The SAP ingress QoS policy number specified on the ingress SAP.
I.MAC/IP	The MAC or IP filter policy ID applied to the ingress SAP.
E.QoS	The SAP egress QoS policy number specified on the egress SAP.
E.Mac/IP	The MAC or IP filter policy ID applied to the egress SAP
A.Pol	The accounting policy ID assigned to the SAP.
Adm	The desired state of the SAP.
Opr	The actual state of the SAP.

Sample Output

```
A:ALA-701# show service sap-using

===============================================================================
Service Access Points
===============================================================================
PortId                 SvcId      Ing.  Ing.    Egr.  Egr.    Anti   Adm  Opr
                                  QoS   Fltr    QoS   Fltr    Spoof
-------------------------------------------------------------------------------
1/1/3                  10203041   1     ip4     1     none    none   Up   Up
1/1/4                  10203042   1     none    1     ip4     none   Up   Up
-------------------------------------------------------------------------------
```

```
Number of SAPs : 2
-------------------------------------------------------------------------------
===============================================================================
A:ALA-701#
```

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, vlan, or vpls.
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.

Label	Description (Continued)
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the Keepalive process.
Oper State	The operational state of the Keepalive process.
Hello Time	Specifies how often the SDP Echo Request messages are transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	Specifies the length of the SDP Echo Request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
E. Fwd. Octets	Specifies the number of forwarded egress octets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. If the SDP type is GRE, then the following message displays: SDP Delivery Mechanism is not MPLS

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 service-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.
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#
```

service-using

Syntax **service-using** [**epipe**] [**ies**] [**vpls**] [**vprn**] [**mirror**] **[apipe] [fpipe] [ipipe]** [**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 **epipe —** Displays matching EPIPE services.

ies — Displays matching IES instances.

vpls — Displays matching VPLS instances.

vprn — Displays matching VPRN services.

mirror — Displays mirror services.

apipe — Displays matching APIPE services.

fpipe — Displays matching FPIPE services.

ipipe — Displays matching IPIPE 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 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-48>show>service# service-using vpls
===============================================================================
Services [vpls]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
```

```
-------------------------------------------------------------------------------
700          VPLS      Up     Down       7                  06/18/2005 09:36:36
725          VPLS      Down   Down       7                  06/18/2005 09:36:36
740          VPLS      Down   Down       1                  06/18/2005 09:36:36
750          VPLS      Down   Down       7                  06/18/2005 09:36:36
1730         VPLS      Down   Down       1730               06/18/2005 09:36:36
9000         VPLS      Up     Down       6                  06/18/2005 09:36:36
90001        VPLS      Up     Down       6                  06/18/2005 09:36:36
-------------------------------------------------------------------------------
Matching Services : 7
-------------------------------------------------------------------------------
===============================================================================
A:ALA-48>show>service#
```

subscriber-using

Syntax **subscriber-using** [**service-id** *service-id*] [**sap-id** *sap-id*] [**interface** *ip-int-name*] [**ip** *ip-address[/mask]*] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*]

Context show>service>subscriber-using

Description Displays subscribers using certain options.

Parameters **service-id *service-id* —** Display subscriber information about the specified service ID.

Values 1 — 2147483647

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port*.*channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda*.*bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094

lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

interface *ip-int-name* — Display subscriber information about the specified interface.

ip *ip-address[/mask* **—** Display subscriber information about the specified IP address.

mac *ieee-address* **—** Display subscriber information about the specified MAC address.

sub-profile *sub-profile-name* **—** Display subscriber information about the specified subscriber profile name.

sla-profile *sla-profile-name* **—** Display subscriber information about the specified SLA profile name.

id

Syntax **id** *service-id*

Context show>service

Description Display 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.

base — Display basic service information.

fdb — Display FDB entries.

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.

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.
Split Horizon Group specifics	
Split Horizon Group	Name of the split horizon group for this VPLS.
Description	Description of the split horizon group .
Last Changed	The date and time of the most recent management-initiated change to this split horizon group.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.
Type	Indicates whether this Service SDP binding is a spoke or a mesh.
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.

Label	Description (Continued)
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the keepalive protocol.
Oper State	The current status of the keepalive protocol.
Hello Time	Specifies how often the SDP Echo Request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP Echo Request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
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 Value	The value of the label used to identify this SAP on the access port.

Label	Description (Continued)
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 of.
Ingress sched-policy	Indicates the ingress QoS scheduler for the SAP.
Egress sched-policy	Indicates the egress QoS scheduler for the SAP.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
SAP Statistics	
Dropped	The number of packets or octets dropped.
Offered Hi Priority	The number of high priority packets, as determined by the SAP ingress QoS policy.
Offered Low Priority	The number of low priority packets, as determined by the SAP ingress QoS policy.
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
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.

Label	Description (Continued)
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Sap per Queue stats	
Ingress Queue 1	The index of the ingress QoS queue of this SAP.
High priority offered	The packets or octets count of the high priority traffic for the SAP.
High priority dropped	The number of high priority traffic packets/octets dropped.
Low priority offered	The packets or octets count of the low priority traffic.
Low priority dropped	The number of low priority traffic packets/octets dropped.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
Out profile forwarded	The number of out-of-profile octets (rate above CIR) forwarded.
Egress Queue 1	The index of the egress QoS queue of the SAP.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
In profile dropped	The number of in-profile packets or octets dropped for the SAP.
Out profile forwarded	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Out profile dropped	The number of out-of-profile packets or octets discarded.
DHCP Relay	
State	Specifies whether DHCP Relay is enabled on this SAP.
Info Option	Specifies whether Option 82 processing is enabled on this SAP.
Action	Specifies the Option 82 processing on this SAP or interface: keep, replace or drop.
Circuit ID	Specifies whether the If Index is inserted in Circuit ID suboption of Option 82.
Remote ID	Specifies whether the far-end MAC address is inserted in Remote ID suboption of Option 82.
STP Service Access Point Specifics	

Label	Description (Continued)
Managed by Service	Specifies the service-id of the management VPLS managing this SAP.
Managed by MSTI	Specifies the MST instance inside the management VPLS managing this SAP.
Last BPDU from	The bridge ID of the sender of the last BPDU received on tbis 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.
Last BPDU from	The bridge ID of the sender of the last BPDU received on this SAP.
Managed by Spoke	Specifies the sap-id inside the management VPLS managing this spoke SDP.
Prune state	Specifies the STP state inherited from the management VPLS.

Sample Output

```
*A:ALA-48# show service id 700 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 700                 Vpn Id            : 0
Service Type      : VPLS
Description       : IMA VPLS
Customer Id       : 7
Last Status Change: 02/17/2007 15:23:16
Last Mgmt Change  : 02/17/2007 15:23:18
Admin State       : Up                  Oper State        : Down
MTU               : 1514                Def. Mesh VC Id   : 700
SAP Count         : 1                   SDP Bind Count    : 2
Send Flush on Fail: Disabled            Host Conn Verify  : Disabled

-------------------------------------------------------------------------------
Split Horizon Group specifics
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Split Horizon Group : DSL-group1
-------------------------------------------------------------------------------
Instance Id       : 1                   Last Change       : 02/17/2007 15:23:18
-------------------------------------------------------------------------------
Split Horizon Group : SHG_test
-------------------------------------------------------------------------------
Description       : test
Instance Id       : 2                   Last Change       : 02/17/2007 15:23:18
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 2:222  -(10.10.10.104)
```

```
-------------------------------------------------------------------------------
Description     : GRE-10.10.10.104
SDP Id            : 2:222                     Type              : Spoke
VC Type           : Ether                     VC Tag            : n/a
Admin Path MTU    : 0                         Oper Path MTU     : 0
Far End           : 10.10.10.104              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
Admin ControlWord : Not Preferred             Oper ControlWord  : False
Last Status Change : 02/17/2007 15:23:16      Signaling         : TLDP
Last Mgmt Change  : 02/17/2007 15:23:18
Flags             : SdpOperDown
                    NoIngVCLabel NoEgrVCLabel
                    PathMTUTooSmall
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

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS


-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move          : Blockable
Stp Admin State   : Up                        Stp Oper State    : Down
Core Connectivity : Down
MCAC Policy Name  :
MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : Down
MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited
Port Role         : N/A                       Port State        : Discarding
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
-------------------------------------------------------------------------------
 Sdp Id 2:700  -(10.10.10.104)
-------------------------------------------------------------------------------
Description     : GRE-10.10.10.104
SDP Id             : 2:700                    Type              : Mesh
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 0
Far End            : 10.10.10.104             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
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Last Status Change : 02/17/2007 15:23:16      Signaling         : TLDP
Last Mgmt Change   : 02/17/2007 15:23:18
Flags              : SdpOperDown
                     NoIngVCLabel NoEgrVCLabel
                     PathMTUTooSmall
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
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

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS


-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/1/9:0
-------------------------------------------------------------------------------
Service Id         : 700
SAP                : 1/1/9:0                  Encap             : q-tag
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100

Admin State        : Up                       Oper State        : Down
Flags              : PortOperDown
Last Status Change : 02/17/2007 15:23:16
Last Mgmt Change   : 02/17/2007 15:23:18
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
```

```
Learned MAC Addr   : 0                        Static MAC Addr   : 0
Admin MTU          : 1518                     Oper MTU          : 1518
Ingress qos-policy : 100                      Egress qos-policy : 1
Shared Q plcy      : default                  Multipoint shared : Enabled
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : 10
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
tod-suite          : None                     qinq-pbit-marking : both
ARP Reply Agent    : Enabled                  Host Conn Verify  : Enabled
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
L2PT Termination   : Disabled                 BPDU Translation  : Disabled

Multi Svc Site     : None
I. Sched Pol       : SLA1
E. Sched Pol       : SLA1
Acct. Pol          : None                     Collect Stats     : Disabled

Anti Spoofing      : None                     Nbr Static Hosts  : 1
Auth Policy        : none
Egr MCast Grp      :
-------------------------------------------------------------------------------
Mac Protection
-------------------------------------------------------------------------------
restrict-protected-src    : Enabled
restrict-unprotected-dest : Disabled

-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Up                       Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Discarding
Port Number        : 2048                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                      Designated Port   : N/A

MCAC Policy Name   :                          MCAC Const Adm St : Enable
MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : Enable
MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited
Forward transitions: 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0
-------------------------------------------------------------------------------
Ingress Queue Override
-------------------------------------------------------------------------------
Queue Id           : 1 (no overrides)
-------------------------------------------------------------------------------
Egress Queue Override
-------------------------------------------------------------------------------
Queue Id           : 1 (no overrides)
-------------------------------------------------------------------------------
DHCP
-------------------------------------------------------------------------------
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-Profile    : None
Def SLA-Profile    : None
Sub-Ident-Policy   : None

Subscriber Limit   : 1
Single-Sub-Parameters
 Prof Traffic Only : False
 Non-Sub-Traffic   : N/A
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
Forwarding Engine Stats
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

Ingress Queue 12 (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 13 (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 15 (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 16 (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 17 (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

-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Down                     Core Connectivity : Down
Stp Admin State    : Up                       Stp Oper State    : Down
Mode               : Rstp                     Vcp Active Prot.  : N/A

Bridge Id          : 10:02.14:30:ff:00:00:00  Bridge Instance Id: 2
Bridge Priority    : 4096                     Tx Hold Count     : 5
Topology Change    : Inactive                 Bridge Hello Time : 5
Last Top. Change   : 0d 00:00:00              Bridge Max Age    : 25
Top. Change Count  : 0                        Bridge Fwd Delay  : 20
MST region revision: 0                        Bridge max hops   : 20
MST region name    :

Root Bridge        : N/A
Primary Bridge     : N/A

Root Path Cost     : 0                        Root Forward Delay: 20
Rcvd Hello Time    : 5                        Root Max Age      : 25
```

```
Root Priority      : 4098                    Root Port         : N/A

-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id        : 700                 Mac Move          : Disabled
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Table Size        : 250                 Total Count       : 1
Learned Count     : 0                   Static Count      : 0
OAM-learned Count : 0                   DHCP-learned Count: 0
host-learned Count: 1
Remote Age        : 900                 Local Age         : 300
High WaterMark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabl               Discard Unknown   : Dsabl
Mac Aging         : Enabl               Relearn Only      : False

-------------------------------------------------------------------------------
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Up
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp                  Oper    MRtr  Send      Max Num   MVR        Num
Id                       State   Port  Queries   Groups    From-VPLS  Groups
-------------------------------------------------------------------------------
sap:1/1/9:0              Down    No    Disabled  No Limit  Local      0
sdp:2:222                Down    No    Disabled  No Limit  N/A        0
sdp:2:700                Down    No    Disabled  No Limit  N/A        0

-------------------------------------------------------------------------------
DHCP Summary, service 700
-------------------------------------------------------------------------------
Sap/Sdp                Snoop  Used/       Arp Reply   Info     Admin
                              Provided    Agent       Option   State
-------------------------------------------------------------------------------
sap:1/1/9:0            No     0/0         Yes         Keep     Down
sdp:2:222              No     N/A         N/A         N/A      N/A
sdp:2:700              No     N/A         N/A         N/A      N/A
-------------------------------------------------------------------------------
Number of Entries : 3
-------------------------------------------------------------------------------
===============================================================================

*A:ALA-48#
```

authentication

Syntax **authentication**

Context show>service>id

Description Enters the context to show session authentication information.

base

Context show>service>id
show>service>id>igmp-snooping

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	The type of service: Epipe, IES, VPLS, VPRN, Mirror, Apipe, Fpipe, Ipipe
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:ALA-48#  show service id 700 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 700                Vpn Id            : 0
Service Type      : VPLS
Description       : IMA VPLS
Customer Id       : 7
Last Status Change: 02/17/2007 15:23:16
Last Mgmt Change  : 02/17/2007 15:23:18
Admin State       : Up                 Oper State        : Down
MTU               : 1514               Def. Mesh VC Id   : 700
SAP Count         : 1                  SDP Bind Count    : 2
Send Flush on Fail: Disabled           Host Conn Verify  : Disabled
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                      Type         AdmMTU  OprMTU  Adm     Opr
-------------------------------------------------------------------------------
sap:1/1/9:0                     q-tag        1518    1518    Up      Down
sdp:2:222 S(10.10.10.104)       n/a          0       0       Up      Down
sdp:2:700 M(10.10.10.104)       n/a          0       0       Up      Down
===============================================================================
*A:ALA-48#
```

dhcp

Syntax **dhcp**

Context show>service>id

Description This command enables the context to display DHCP information for the specified service.

Sample Output

```
A:ALA-48>show>service>id>dhcp#
-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Mac Move          : Blockable
Stp Admin State   : Up                       Stp Oper State    : Down
Core Connectivity : Down
Port Role         : N/A                      Port State        : Discarding
Port Number       : 2048                     Port Priority     : 128
Port Path Cost    : 10                       Auto Edge         : Enabled
Admin Edge        : Disabled                 Oper Edge         : N/A
Link Type         : Pt-pt                    BPDU Encap        : Dot1d
Root Guard        : Disabled                 Active Protocol   : N/A
Last BPDU from    : N/A
CIST Desig Bridge : N/A                      Designated Port   : N/A

Forward transitions: 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd    : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd    : 0                        TCN BPDUs tx      : 0
```

```
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0
-------------------------------------------------------------------------------
DHCP
-------------------------------------------------------------------------------
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-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-48>show>service>id>dhcp#
```

lease-state

Syntax **lease-state** [[**sap** *sap-id*] [**sdp** [*sdp-id*[:*vc-id*]]] | [**interface** *interface-name*] | [**ip-address** *ip-address[/mask*]] | [**mac** *ieee-address*] | [**wholesaler** *service-id*]] [**detail**]

Context show>service>id>dhcp

Description This command displays DHCP lease state information. Note that the **wholesaler** *service-id* parameter is applicable only in the VPRN context.

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword

	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp *sdp-id* — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Values 1 — 4294967295

interface *interface-name* — Displays information for the specified IP interface.

ip-address *ip-address* — Displays information associated with the specified IP address.

detail — Displays detailed information.

wholesaler *service-id* **—** The VPRN service ID of the wholesaler. When specified in this context, SAP, SDP, interface, IP address and MAC parameters are ignored.

Values 1 — 2147483647

Sample Output

```
A:ALA-Dut-A# show service id 13 dhcp lease-state
===============================================================================
DHCP lease state table, service 13
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
13.13.40.1      00:00:00:00:00:13 1/1/1:13            00h00m58s  Radius
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-Dut-A#


A:subscr_mgt# show service id 1 dhcp lease-state detail
===============================================================================
DHCP lease states for service 1
===============================================================================
```

```
-------------------------------------------------------------------------------
DHCP Lease 1.1.1.2
-------------------------------------------------------------------------------
Mac Address          : 00:01:00:00:00:01
SAP                  : 1/1/1:1
Remaining Lifetime   : 03h28m19s
Persistence Key      : 0x00000000

Sub-Ident            : "alcatel_A_1"
Sub-Profile-String   : "sub_prof_B_2"
SLA-Profile-String   : "sla_prof_C_3"
-------------------------------------------------------------------------------
DHCP Lease 1.1.1.3
-------------------------------------------------------------------------------
Mac Address          : 00:01:00:00:00:02
SAP                  : 1/1/1:1
Remaining Lifetime   : 18h17m12s
Persistence Key      : 0x00000001

Sub-Ident            : "alcatel_A_1"
Sub-Profile-String   : "sub_prof_B_2"
SLA-Profile-String   : "sla_prof_C_3"
-------------------------------------------------------------------------------
DHCP Lease 1.1.1.4
-------------------------------------------------------------------------------
Mac Address          : 00:01:00:00:00:03
SAP                  : 1/1/1:1
Remaining Lifetime   : 01d09h06m
Persistence Key      : 0x00000002

Sub-Ident            : "alcatel_A_1"
Sub-Profile-String   : "sub_prof_B_2"
SLA-Profile-String   : "sla_prof_C_3"
-------------------------------------------------------------------------------
Number of lease states : 3
===============================================================================
A:subscr_mgt#
```

statistics

Syntax **statistics** [**sap** *sap-id*
statistics [**sdp** *sdp-id*:*vc-id*]
statistics [**interface** *interface-name*]

Context show>service>id>dhcp

Description Displays DHCP statistics information.

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

Values T*sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*
port-id	*slot/mda/port*[.*channel*]

aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp *sdp-id* — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Values 1 — 4294967295

interface *interface-name* — Displays information for the specified IP interface.

Show DHCP Statistics Output — The following table describes the output fields for DHCP statistics.

Label	Description
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped.
Client Packets Forwarded	The number of packets received from the DHCP clients that were forwarded.
Client Packets Dropped	The number of packets received from the DHCP clients that were dropped.

Label	Description
Server Packets Snooped	The number of packets received from the DHCP server that were snooped.
Server Packets Forwarded	The number of packets received from the DHCP server that were forwarded.
Server Packets Dropped	The number of packets received from the DHCP server that were dropped.

Sample Output

```
*A:7450# show service id 1 dhcp statistics
=====================================================
DHCP Statistics, service 1
=====================================================
Client Packets Snooped               : 0
Client Packets Forwarded             : 0
Client Packets Dropped               : 0
Client Packets Proxied (RADIUS)      : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Snooped               : 0
Server Packets Forwarded             : 0
Server Packets Dropped               : 0
DHCP RELEASEs Spoofed                : 0
DHCP FORCERENEWs Spoofed             : 0
=====================================================
A:test#
```

summary

Syntax **summary**

Context show>service>id>dhcp

Description **Show Service-ID DHCP Summary —** The following table describes show service-id DHCP summary output fields:

Label	Description
Sap/Sdp	The configuration identification, expressed by a string containing “card/mda/port/:logical-id”.
Snoop	Yes — The packets received from the DHCP clients were snooped.
	No - The packets received from the DHCP clients were not snooped
Used/Provided	Used — The number of lease-states that are currently in use on a specific interface, that is, the number of clients on that interface got an IP address by DHCP. This value is always less than or equal to the ‘Provided’ field.
	Provided — The lease-populate value that is configured for a specific interface.

Label	Description
Arp Reply Agent	Displays whether or not there is proper handling of received ARP requests from subscribers.
Info Option	Keep – The existing information is kept on the packet and the router does not add any additional information.
	Replace – On ingress, the existing information-option is replaced with the information-option from the router.
	Drop – The packet is dropped and an error is logged.
Admin State	Indicates the administrative state.

```
A:test>config>service# show service id 100 dhcp summary
===============================================================================
DHCP Summary, service 100
===============================================================================
Sap/Sdp                 Snoop  Used/       Arp Reply   Info     Admin
                               Provided    Agent       Option   State
-------------------------------------------------------------------------------
sap:1/2/1:0.*           No     0/0         No          Keep     Down
sap:3/1/2:0.*           No     0/0         No          Keep     Down
sap:10/1/1.4.3.2.24:1* No      0/0         No          Keep     Down
sap:lag-1:100           No     0/0         No          Keep     Down
sdp:11:100              No     N/A         N/A         N/A      N/A
sdp:20:100              No     N/A         N/A         N/A      N/A
sdp:30:100              No     N/A         N/A         N/A      N/A
sdp:500:100             No     N/A         N/A         N/A      N/A
-------------------------------------------------------------------------------
Number of Entries : 8
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
A:test>config>service#
```

fdb

Syntax **fdb sap** *sap-id* [**expiry**]] | [**sdp** *sdp-id* [**expiry**]] | [**mac** *ieee-address* [**expiry**]] | [**detail**][**expiry**]

Context show>service>id
show>service>fdb-mac

Description Display FDB entries for a given MAC address.

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

Sample Output

```
A:ALA-48>show>service>id# fdb mac
===============================================================================
Service Forwarding Database
===============================================================================
ServId    MAC               Source-Identifier     Type/Age  Last Change
-------------------------------------------------------------------------------
6         00:aa:00:00:00:00 sap:lag-2             L/0       06/27/2006
15:04:31
6         00:aa:00:00:00:01 sap:lag-2             L/0       06/27/2006
15:04:31
6         00:aa:00:00:00:02 sap:lag-2             L/0       06/27/2006
15:04:31
6         00:aa:00:00:00:03 sap:lag-2             L/0       06/27/2006
15:04:31
6         00:aa:00:00:00:04 sap:lag-2             L/0       06/27/2006
15:04:31
10        12:12:12:12:12:12 sap:1/1/1:100         S         06/26/2006
10:03:29
===============================================================================
A:ALA-48>show>service>id#


A:ALA-48# show service id 700 fdb [mac-protect]
===============================================================================
Forwarding Database, Service <service-id>
```

```
===============================================================================
ServId    MAC               Source-Identifier     Type/Age  Last Change
-------------------------------------------------------------------------------
1         aa:aa:aa:aa:aa:aa sdp:100:1             P         11/02/2006 06:04:03
-------------------------------------------------------------------------------
No. of MAC Entries: 1
===============================================================================
A:ALA-48#
```

egress-multicast-group

Syntax **egress-multicast-group** [*group-name*]

Context show>service

Description Displays egress multicast group information.

Parameters *group-name* — Specifies the name of the egress multicast group.

Sample Output

```
A:Dut-C# show service egress-multicast-group emg1
===============================================================================
Egress Multicast Group Entry
===============================================================================
Group              : emg1
-------------------------------------------------------------------------------
Chain Limit        : 16                Encap Type         : dot1q
Dot1q ether type   : 0x8100            Filter-Id          : n/a
-------------------------------------------------------------------------------
Service Access Points
1/1/1:100
-------------------------------------------------------------------------------
===============================================================================
A:Dut-C#
```

gsmp

Syntax **gsmp**

Context show>service>id

Description This command displays 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                           Neigbor           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                           Neigbor           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                           Neigbor           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.

Description **Show Sessions Neighbor Output —** The following table describes show sessions neighbor output fields:

Label	Description
State	The current state of the ANCP session.
Peer Instance	The instance number of the ANCP session at the neighbor's side.
Sender Instance	The instance number of the ANCP session at our side.
Peer Port	The port number of the ANCP session at the neighbor's side.
Sender Port	The port number of the ANCP session at the local side.
Peer Name	The MAC address of the ANCP session at the neighbor's side.
Sender name	The mac address of the ANCP session at the local side.
timeouts	The number of adjacency protocol message timeouts.
Max. Timeouts	The maximum allowed of the above timeouts before closing.
Peer Timer	The timer value for the neighbor peridodic adjacency protocol messages.
Sender Timer	The timer value for the local peridodic adjacency protocol messages.
Capabilities	The negotiated capabilties for the Established ANCP session (DTD : dynamic topology discovery - OAM : operation and maintenance).
Conf Capabili-ties	The configured local capabilities.
Priority Mark-ing	The DSCP bits for the IP messages used in the ANCP session.
Local Addr.	The destination IP address for this ANCP session.

Label	Description
Conf Local Addr.	The destination IP address accepted for ANCP connections.

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), neigbor 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), neigbor 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
===============================================================================
ANCP-String                                                    Assoc.
State
-------------------------------------------------------------------------------
7330-ISAM-E47 atm 1/1/01/01:19425.64048                        ANCP    Up
-------------------------------------------------------------------------------
Number of ANCP-Strings : 1
===============================================================================
```

host

Syntax **host** [**sap** *sap-id*] [**detail**]
host summary

Context show>service>id

Description This command displays static host information configured on this service.

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword

	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

summary — Displays summary host information.

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* — Specifies the physical port identifier portion of the SAP definition.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1.vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*
port-id	*slot/mda/port*[.*channel*]
aps-id	aps-*group-id*[.*channel*]
	aps keyword
	group-id 1 — 64
bundle-*type-slot/mda.bundle-num*	
	bundle keyword
	type ima, ppp
	bundle-num 1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*
	bpgrp keyword
	type ima
	bpgrp-num 1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*
	ccag keyword
	id 1 — 8
	path-id a, b
	cc-type .sap-net, .net-sap]

	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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 1/2/3:0143.144.145.1                                 Up
===============================================================================
A:ALA-48>show>service>id#
```

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#
```

mfib

Syntax **mfib brief**
mfib [**group** *grp-address*]
mfib statistics [**group** *grp-address*]

Context show>service>id>

Description Displays the multicast FIB on the VPLS service.

Parameters **group grp-ip-address** — Displays the multicast FIB for a specific multicast group address.

statistics — Displays statistics on the multicast FIB.

Output **Show Output —** The following table describes the command output fields:

Label	Description
Source Address	IPv4 unicast source addres

Label	Description (Continued)
Group Address	IPv4 multicast group address.
SAP/SDP ID	Indicates the SAP/SDP to which the corresponding multicast stream will be forwarded/blocked.
Forwarding/Block-ing	Indicates whether the corresponding multicast stream will be blocked/forwarded.
Number of Entries	Number of entries in the MFIB.
Forwarded Packets	Indicates the number of multicast packets forwarded for the corresponding source/group.
Forwarded Octets	Indicates the number of octets forwarded for the corresponding source/group.

Sample Output

```
*A:ALA-1>show>service>id # mfib
===============================================================================
IGMP Snooping MFIB for service 10
===============================================================================
Source Address  Group Address   Sap/Sdp Id             Fwd/Blk
-------------------------------------------------------------------------------
*               225.0.0.1       sap:2/1/5:1            Fwd
*               225.0.0.7       sap:2/1/5:7            Fwd
-------------------------------------------------------------------------------
Number of entries: 2

*A:ALA-SR12-D#  show service id 1 mfib mesh-sdp 41:1 igmp-snooping no mrouter-port
===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address       Sap/Sdp Id               Svc Id   Fwd/Blk
-------------------------------------------------------------------------------
*               *                       sap:6/1/1:1              Local    Fwd
                                  (!) sdp:41:3                 Local    Fwd
                                      sdp:41:5                 Local    Fwd
*               225.0.0.0             sap:6/1/1:1              Local    Fwd
                                  (!) sdp:41:1                 Local    Fwd
                                  (!) sdp:41:3                 Local    Fwd
                                      sdp:41:5                 Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
(!) Allowed-MDA-destination restrictions apply
*A:ALA-SR12-D#



*A:ALA-1>show>service>id # mfib statistics
===============================================================================
IGMP Snooping MFIB for service 10
===============================================================================
Source Address  Group Address   Fwd Pkts              Fwd Octets
-------------------------------------------------------------------------------
1.1.1.1         225.0.0.1       291                    9281
```

```
1.1.1.2         225.0.0.1       0                       0
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
```

mstp-configuration

Syntax **mstp-configuration**

Context show>service>id

Description This command displays the MSTP specific configuration data. This command is only valid on a management VPLS.

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.

wholesalers

Syntax **wholesalers**

Context show>service>id

Description This command displays service wholesaler information.

sap

Syntax **sap** *sap-id* [**detail**]

Context show>service>id

Description This command displays information for the SAPs associated with the service.

If no optional parameters are specified, a summary of all associated SAPs is displayed.

Parameters **sap** *sap-id* — The ID that displays SAPs for the service in the *slot*/*mda*/*port*[**.***channel*] form.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*

cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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 SAP and qtag.
Encap	The encapsulation type of the SAP.
Ethertype	Specifies an Ethernet type II Ethertype value.
Admin State	The administrative state of the SAP.
Oper State	The operating state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP.
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.

Label	Description
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.
Forwarding Engine Stats	
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.
Off. Uncolor	The number of uncolored packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Queueing Stats (Ingress QoS Policy)	
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.
Queuing Stats (Egress QoS Policy)	
Dro. InProf	The number of in-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.

Label	Description
Dro. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, 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 Han-dling	The indication that OAM cells are being processed.
AAL-5 Encap	The AAL-5 encapsulation type.

Sample Output

```
A:ALA-48>show>service>id# sap 1/1/9:0
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 700
SAP                : 1/1/9:0                  Encap             : q-tag
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100

Admin State        : Up                       Oper State        : Down
Flags              : PortOperDown
Last Status Change : 02/17/2007 15:23:16
Last Mgmt Change   : 02/17/2007 15:23:18
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
Admin MTU          : 1518                     Oper MTU          : 1518
Ingress qos-policy : 100                      Egress qos-policy : 1
Shared Q plcy      : default                  Multipoint shared : Enabled
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : 10
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
ARP Reply Agent    : Enabled                  Host Conn Verify  : Enabled
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
L2PT Termination   : Disabled                 BPDU Translation  : Disabled

Multi Svc Site     : None
I. Sched Pol       : SLA1
E. Sched Pol       : SLA1
Acct. Pol          : None                     Collect Stats     : Disabled

Anti Spoofing      : None                     Nbr Static Hosts  : 1
MCAC Policy Name   :                          MCAC Const Adm St : Enable
MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited
```

```
-------------------------------------------------------------------------------
Ingress Queue Override
-------------------------------------------------------------------------------
Queue Id           : 1 (no overrides)
-------------------------------------------------------------------------------
Egress Queue Override
-------------------------------------------------------------------------------
Queue Id           : 1 (no overrides)
-------------------------------------------------------------------------------
DHCP
-------------------------------------------------------------------------------
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-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-48>show>service>id#


A:ALA-48>show>service>id# sap 1/1/9:0 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 700
SAP                : 1/1/9:0                 Encap             : q-tag
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100

Admin State        : Up                      Oper State        : Down
Flags              : PortOperDown
Last Status Change : 06/07/2006 12:30:02
Last Mgmt Change   : 06/07/2006 12:30:03
Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
Admin MTU          : 1518                    Oper MTU          : 1518
Ingress qos-policy : 100                     Egress qos-policy : 1
Shared Q plcy      : default                 Multipoint shared : Disabled
Ingress Filter-Id  : n/a                     Egress Filter-Id  : n/a
tod-suite          : None
ARP Reply Agent    : Disabled
Mac Learning       : Enabled                 Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                 Mac Pinning       : Disabled
L2PT Termination   : Disabled                BPDU Translation  : Disabled

Multi Svc Site     : None
I. Sched Pol       : SLA1
E. Sched Pol       : SLA1
Acct. Pol          : None                    Collect Stats     : Disabled

Anti Spoofing      : None                    Nbr Static Hosts  : 0
```

```
Auth Policy        : none
Egr MCast Grp      :
-------------------------------------------------------------------------------
Mac Protection
-------------------------------------------------------------------------------
restrict-protected-src    : Disabled
restrict-unprotected-dest : Disabled

-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Up                       Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Discarding
Port Number        : 2048                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                      Designated Port   : N/A

MCAC Policy Name   :                          MCAC Const Adm St : Enable
MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : Enable
MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited
Forward transitions: 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0
-------------------------------------------------------------------------------
DHCP
-------------------------------------------------------------------------------
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-Profile    : None
Def SLA-Profile    : None
Sub-Ident-Policy   : None

Subscriber Limit   : 1
Single-Sub-Parameters
 Prof Traffic Only : False
 Non-Sub-Traffic   : N/A
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                      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 23 (Multipoint) (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 25 (Multipoint) (Priority)
Off. HiPrio *A:ALA-48# show service id 700 sap 1/1/9:0 detail

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 700
SAP                : 1/1/9:0                  Encap             : q-tag
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100

Admin State        : Up                       Oper State        : Down
Flags              : PortOperDown
Last Status Change : 02/17/2007 15:23:16
Last Mgmt Change   : 02/17/2007 15:23:18
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
Admin MTU          : 1518                     Oper MTU          : 1518
Ingress qos-policy : 100                      Egress qos-policy : 1
Shared Q plcy      : default                  Multipoint shared : Enabled
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : 10
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
ARP Reply Agent    : Enabled                  Host Conn Verify  : Enabled
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
L2PT Termination   : Disabled                 BPDU Translation  : Disabled

Multi Svc Site     : None
I. Sched Pol       : SLA1
```

```
E. Sched Pol       : SLA1
Acct. Pol          : None                     Collect Stats     : Disabled

Anti Spoofing      : None                     Nbr Static Hosts  : 1
MCAC Policy Name   :                          MCAC Const Adm St : Enable
MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited
Auth Policy        : none
Egr MCast Grp      :
-------------------------------------------------------------------------------
Mac Protection
-------------------------------------------------------------------------------
restrict-protected-src    : Enabled
restrict-unprotected-dest : Disabled

-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Up                       Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Discarding
Port Number        : 2048                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                      Designated Port   : N/A

Forward transitions: 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0
-------------------------------------------------------------------------------
Ingress Queue Override
-------------------------------------------------------------------------------
Queue Id           : 1 (no overrides)
-------------------------------------------------------------------------------
Egress Queue Override
-------------------------------------------------------------------------------
Queue Id           : 1 (no overrides)
-------------------------------------------------------------------------------
DHCP
-------------------------------------------------------------------------------
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-Profile    : None
Def SLA-Profile    : None
Sub-Ident-Policy   : None

Subscriber Limit   : 1
Single-Sub-Parameters
```

```
 Prof Traffic Only : False
 Non-Sub-Traffic   : N/A
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
Forwarding Engine Stats
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

Ingress Queue 12 (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 13 (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 15 (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 16 (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 17 (Unicast) (Priority)
Off. HiPrio           : 0                      0
Off. LoPrio           : 0                      0
Dro. HiPrio           : 0                      0
Dro. LoPrio           : 0                      0
For. InProf           : 0                      0
For. OutProf          : 0                      0

Egress Queue 1
For. InProf           : 0                      0
For. OutProf          : 0                      0
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0
===============================================================================
A:ALA-48>show>service>id#
```

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.

Label	Description (Continued)
Split Horizon Group	Name of the split horizon group that the SDP belongs to.
VC Type	Displays the VC type: ether, vlan, or vpls.
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 current status of the SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the Keepalive process.
Oper State	The operational state of the Keepalive process.
Hello Time	Specifies how often the SDP Echo Request messages are transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	Specifies the length of the SDP Echo Request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.

Label	Description (Continued)
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
E. Fwd. Octets	Specifies the number of forwarded egress octets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. If the SDP type is GRE, then the following message displays: SDP Delivery Mechanism is not MPLS

Sample Output

```
*A:ALA-12#  show service id 9000 sdp 2:22 detail
===============================================================================
Service Destination Point (Sdp Id : 2:22) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 2:22  -(10.10.10.103)
-------------------------------------------------------------------------------
Description     : GRE-10.10.10.103
SDP Id            : 2:22                      Type              : Spoke
Split Horiz Grp : (DSL-group1
VC Type           : Ether                     VC Tag            : n/a
Admin Path MTU    : 4462                      Oper Path MTU     : 4462
Far End           : 10.10.10.103              Delivery          : GRE
Admin State       : Up                        Oper State        : TLDP Down
Ingress Label     : 0                         Egress Label      : 0
Ingress Filter    : n/a                       Egress Filter     : n/a
Last Changed      : 10/29/2006 11:48:20       Signaling         : TLDP

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
E. Fwd. Pkts.     : 0                         E. Fwd. Octets    : 0

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS


-------------------------------------------------------------------------------
Rstp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move          : Disabled
Rstp Admin State  : Up                        Rstp Oper State   : Down
Core Connectivity : Down
Port Role         : N/A                       Port State        : Discarding
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
Designated Bridge  : N/A                     Designated Port Id: 0
Active Protocol    : N/A

Fwd Transitions    : 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALA-12#
```

split-horizon-group

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

Context show>service>id

Description This command displays service split horizon groups.

```
*A:ALA-1# show service id 700 split-horizon-group
===============================================================================
Service: Split Horizon Group
===============================================================================
Name                            Description
-------------------------------------------------------------------------------
DSL-group1                      Split horizon group for DSL
-------------------------------------------------------------------------------
No. of Split Horizon Groups: 1
===============================================================================
*A:ALA-1#

*A:ALA-1# show service id 700 split-horizon-group DSL-group1
===============================================================================
Service: Split Horizon Group
===============================================================================
Name                            Description
-------------------------------------------------------------------------------
DSL-group1                      Split horizon group for DSL
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
SAP                             1/1/3:1

SDP                             108:1
SDP                             109:1

-------------------------------------------------------------------------------
SAPs Associated : 1            SDPs Associated : 2
===============================================================================
*A:ALA-1#
```

stp

Syntax **stp** [**detail**]

Context show>service>id

Description Displays information for the spanning tree protocol instance for the service.

Parameters **detail** — Displays detailed information.

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

Label	Description
RSTP Admin State	Indicates the administrative state of the Rapid Spanning Tree Protocol instance associated with this service.
Core Connectivity	Indicates the connectivity status to the core.
RSTP Oper State	Indicates the operational state of the Rapid Spanning Tree Protocol instance associated with this service. This field is applicable only when STP is enabled on the router.
Bridge-id	Specifies the MAC address used to identify this bridge in the network.
Hold Time	Specifies the interval length during which no more than two Configuration BPDUs shall be transmitted by this bridge.
Bridge fwd delay	Specifies how fast a bridge changes its state when moving toward the forwarding state.
Bridge Hello time	Specifies the amount of time between the transmission of Configuration BPDUs.
Bridge max age	Specifies the maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded. This is the actual value that this bridge is currently using.
Bridge priority	Defines the priority of the Spanning Tree Protocol instance associated with this service.
Topology change	Specifies whether a topology change is currently in progress.
Last Top. change	Specifies the time (in hundredths of a second) since the last time a topology change was detected by the Spanning Tree Protocol instance associated with this service.
Top. change count	Specifies the total number of topology changes detected by the Spanning Tree Protocol instance associated with this service since the management entity was last reset or initialized.
Root bridge-id	Specifies the bridge identifier of the root of the spanning tree as determined by the Spanning Tree Protocol instance associated with this service. This value is used as the Root Identifier parameter in all Configuration BPDUs originated by this node.
Root path cost	Specifies the cost of the path to the root bridge as seen from this bridge.

Label	Description
Root forward delay	Specifies how fast the root changes its state when moving toward the forwarding state.
Root hello time	Specifies the amount of time between the transmission of configuration BPDUs.
Root max age	Specifies the maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded.
Root priority	This object specifies the priority of the bridge that is currently selected as root-bridge for the network.
Root port	Specifies the port number of the port which offers the lowest cost path from this bridge to the root bridge.
SAP Identifier	The ID of the access port where this SAP is defined.
RSTP State	The operational state of RSTP.
STP Port State	Specifies the port identifier of the port on the designated bridge for this port's segment.
BPDU encap	Specifies the type of encapsulation used on BPDUs sent out and received on this SAP.
Port Number	Specifies the value of the port number field which is contained in the least significant 12 bits of the 16-bit Port ID associated with this SAP.
Priority	Specifies the value of the port priority field which is contained in the most significant 4 bits of the 16-bit Port ID associated with this SAP.
Cost	Specifies the contribution of this port to the path cost of paths towards the spanning tree root which include this port.
Fast Start	Specifies whether Fast Start is enabled on this SAP.
Designated Port	Specifies the port identifier of the port on the designated bridge for this port's segment.
Designated Bridge	Specifies the bridge identifier of the bridge which this port considers to be the designated bridge for this port's segment.
Service Access Points	
Managed by Service	Specifies the service-id of the management VPLS managing this SAP or spoke SDP.
Managed by SAP/ spoke	Specifies the sap-id or sdp-id inside the management VPLS managing this SAP or spoke SDP.
Prune state	Specifies the STP state inherited from the management VPLS.

Sample Output

```
*A:ALA-12# show service id 11 stp
===============================================================================
Stp info, Service 11
===============================================================================
Bridge Id          : 80:00.22:68:ff:00:00:00  Top. Change Count : 1
Root Bridge        : 00:00.22:69:ff:00:00:00  Stp Oper State    : Syncing Vcp
Primary Bridge     : N/A                      Topology Change   : Inactive
Mode               : Mstp                     Last Top. Change  : 0d 19:12:58
Vcp Active Prot.   : N/A
Root Port          : 2048                     External RPC      : 10

===============================================================================
MSTP specific info for CIST
===============================================================================
Regional Root      : This Bridge              Root Port         : 2048
Internal RPC       : 0                        Remaining Hopcount: 20

===============================================================================
Stp port info for CIST
===============================================================================
Sap/Sdp Id         Oper-     Port-      Port-       Port-  Oper-  Link-  Active
                   State     Role       State       Num    Edge   Type   Prot.
-------------------------------------------------------------------------------
1/1/1:0            Up        Root       Forward     2048   False  Pt-pt  Mstp
1/1/3:0            Up        N/A        Forward     2049   N/A    Pt-pt  N/A
1/1/4:*            Up        Designated Forward     2050   False  Pt-pt  Mstp
===============================================================================

===============================================================================
MSTP specific info for MSTI 111
===============================================================================
Regional Root      : 80:6f.1c:65:ff:00:00:00  Root Port         : 2050
Internal RPC       : 10                       Remaining Hopcount: 19

===================================================================
MSTP port info for MSTI 111
===================================================================
Sap/Sdp Id         Oper-     Port-      Port-       Port-  Same
                   State     Role       State       Num    Region
-------------------------------------------------------------------
1/1/1:0            Up        Master     Forward     2048   False
1/1/3:0            Up        N/A        Forward     2049   N/A
1/1/4:*            Up        Root       Forward     2050   True
===================================================================
*A:ALA-12#


*A:ALA-12# show service id stp detail
===============================================================================
Spanning Tree Information
===============================================================================
-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up                       Core Connectivity : Down
Stp Admin State    : Up                       Stp Oper State    : Up
Mode               : Mstp                     Vcp Active Prot.  : N/A

Bridge Id          : 80:00.22:68:ff:00:00:00  Bridge Instance Id: 0
Bridge Priority    : 32768                    Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 0d 19:14:34              Bridge Max Age    : 20
Top. Change Count  : 1                        Bridge Fwd Delay  : 15
```

```
MST region revision: 0                        Bridge max hops   : 20
MST region name    : abc

Root Bridge        : 00:00.22:69:ff:00:00:00
Primary Bridge     : N/A

Root Path Cost     : 10                       Root Forward Delay: 15
Rcvd Hello Time    : 2                        Root Max Age      : 20
Root Priority      : 0                        Root Port         : 2048

MSTP info for CIST :
Regional Root      : This Bridge              Root Port         : 2048
Internal RPC       : 0                        Remaining Hopcount: 20
MSTP info for MSTI 111 :
Regional Root      : 80:6f.1c:65:ff:00:00:00  Root Port         : 2050
Internal RPC       : 10                       Remaining Hopcount: 19

-------------------------------------------------------------------------------
Spanning Tree Virtual Core Port (VCP) Specifics
-------------------------------------------------------------------------------
Mesh Sdp Id        Sdp           Sdp Bind     Mesh Sdp     HoldDown   Awaiting
                   Oper-state    Oper-state   Port-state   Timer      Agreement
-------------------------------------------------------------------------------
3:11               Down          Down         Discard      Inactive   N/A
4:11               Down          Down         Discard      Inactive   N/A
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Spanning Tree Sap/Spoke SDP Specifics
-------------------------------------------------------------------------------
SAP Identifier     : 1/1/1:0                  Stp Admin State   : Up
Port Role          : Root                     Port State        : Forwarding
Port Number        : 2048                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Mstp
Last BPDU from     : 00:00.22:69:ff:00:00:00  Inside Mst Region : False
CIST Desig Bridge  : 00:00.22:69:ff:00:00:00  Designated Port   : 34816
MSTI 111 Port Prio : 128                      Port Path Cost    : 10
MSTI 111 Desig Brid: This Bridge              Designated Port   : 34816
Forward transitions: 1                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 34638                    MST BPDUs tx      : 3

SAP Identifier     : 1/1/3:0                  Stp Admin State   : Down
Port Role          : N/A                      Port State        : Forwarding
Port Number        : 2049                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                      Designated Port   : 0
MSTI 111 Port Prio : 128                      Port Path Cost    : 10
MSTI 111 Desig Brid: N/A                      Designated Port   : 0
Forward transitions: 1                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0
```

```
SAP Identifier     : 1/1/4:*                  Stp Admin State   : Up
Port Role          : Designated               Port State        : Forwarding
Port Number        : 2050                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Mstp
Last BPDU from     : 50:00.1c:65:ff:00:00:00  Inside Mst Region : True
CIST Desig Bridge  : This Bridge              Designated Port   : 34818
MSTI 111 Port Prio : 128                      Port Path Cost    : 10
MSTI 111 Desig Brid: 80:6f.1c:65:ff:00:00:00  Designated Port   : 34819
Forward transitions: 1                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 34636                    MST BPDUs tx      : 34640

===============================================================================
```

subscriber-hosts

Syntax **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-address*[/*mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**detail**]

Context show>service>id

Description Displays subscriber host information.

Parameters **sap** *sap-id* — Displays the specified subscriber host SAP information.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword

	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

ip-address/mask — 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.

Values Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /30 subnets).
mask: 1 — 32

mac *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.

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

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

detail — Displays detailed information.

Sample Output

```
A:A:ALA#-SR12# show service id 20 subscriber-hosts
===============================================================================
Subscriber Host table
===============================================================================
Sap Id                 IP Address      MAC Address       Origin(*) Subscriber
-------------------------------------------------------------------------------
1/2/6:0                101.1.1.10      00:bb:bb:00:00:00 S/-/-
   Elisa-20-static
-------------------------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================================
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
A:ALA#

A:ALA# show service id 10 subscriber-hosts
===============================================================================
Subscriber Host table
===============================================================================
```

```
Sap Id                 IP Address      MAC Address       Origin(*) Subscriber
-------------------------------------------------------------------------------
1/2/5:0                100.1.1.10      00:aa:aa:00:00:01 -/D/-
   SUB-10-00aaaa000001
-------------------------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================================
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
A:ALA-SR12#
```

statistics

Syntax **statistics** [**policy** *name*] [**sap** *sap-id*]

Context show>service>id>authentication

Description Displays session authentication statistics for this service.

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type*-*slot/mda*.*bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type*-*bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255

vci	1, 2, 5 — 65535
dlci	16 — 1022

Sample Output

```
*A:ALA-1# show service id 11 authentication statistics
===============================================================
Authentication statistics
===============================================================
Interface / SAP                Authentication  Authentication
                               Successful      Failed
---------------------------------------------------------------
vpls-11-90.1.0.254              1582            3
---------------------------------------------------------------
Number of entries: 1
===============================================================
*A:ALA-1#
```

IGMP Snooping Show Commands

igmp-snooping

Syntax **igmp-snooping**

Context show>service>id>

Description This command enables the context to display IGMP snooping information.

all

Syntax **all**

Context show>service>id>igmp-snooping

Description Displays detailed information for all aspects of IGMP snooping on the VPLS service.

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

Label	Description
Admin State	The administrative state of the IGMP instance.
Querier	Displays the address of the IGMP querier on the IP subnet to which the interface is attached.
Sap/Sdp Id	Displays the SAP and SDP IDs of the service ID.
Oper State	Displays the operational state of the SAP and SDP IDs of the service ID.
Mrtr Port	Specifies if the port is a multicast router port.
Send Queries	Specifies whether the send-queries command is enabled or disabled.
Max Num Groups	Specifies the maximum number of multicast groups that can be joined on this SAP or SDP.
MVR From VPLS	Specifies MVR from VPLS.
Num Groups	Specifies the actual number of multicast groups that can be joined on this SAP or SDP.

Sample Output

```
A:ALA-48>show>service>id>igmp-snooping>snooping# all
===============================================================================
IGMP Snooping info for service 750
===============================================================================
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Up
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp                Oper      MRtr  Send      Max Num   Num
Id                     State     Port  Queries   Groups    Groups
-------------------------------------------------------------------------------
sap:1/1/7:0            Down      No    Disabled  No Limit  0
sdp:1:22               Down      No    Disabled  No Limit  0
sdp:8:750              Down      No    Disabled  No Limit  0
-------------------------------------------------------------------------------
IGMP Snooping Querier info
-------------------------------------------------------------------------------
No querier found for this service.
-------------------------------------------------------------------------------
IGMP Snooping Multicast Routers
-------------------------------------------------------------------------------
MRouter          Sap/Sdp Id             Up Time        Expires   Version
-------------------------------------------------------------------------------
Number of mrouters: 0
-------------------------------------------------------------------------------
IGMP Snooping Proxy-reporting DB
-------------------------------------------------------------------------------
Group Address    Mode     Type     Up Time        Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping SAP 1/1/7:0 Port-DB
-------------------------------------------------------------------------------
Group Address    Mode     Type     Up Time        Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping SDP 1:22 Port-DB
-------------------------------------------------------------------------------
Group Address    Mode     Type     Up Time        Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping SDP 8:750 Port-DB
-------------------------------------------------------------------------------
Group Address    Mode     Type     Up Time        Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups
-------------------------------------------------------------------------------
IGMP Snooping Statistics
-------------------------------------------------------------------------------
Message Type            Received      Transmitted      Forwarded
-------------------------------------------------------------------------------
General Queries         0             0                0
Group Queries           0             0                0
```

```
Group-Source Queries    0             0               0
V1 Reports              0             0               0
V2 Reports              0             0               0
V3 Reports              0             0               0
V2 Leaves               0             0               0
Unknown Type            0             N/A             0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length              : 0
Bad IP Checksum         : 0
Bad IGMP Checksum       : 0
Bad Encoding            : 0
No Router Alert         : 0
Zero Source IP          : 0

Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0
===============================================================================
A:ALA-48>show>service>id>snooping#
```

mrouters

Syntax **mrouters** [**detail**]

Context show>service>id>igmp-snooping

Description Displays all multicast routers.

Parameters **detail —** Displays detailed information.

```
*A:rbae_C# show service id 1 igmp-snooping mrouters
===============================================================================
IGMP Snooping Multicast Routers for service 1
===============================================================================
MRouter          Sap/Sdp Id             Up Time        Expires   Version
-------------------------------------------------------------------------------
10.10.1.1        2/1/5:1                0d 00:00:26    14s       3
10.20.1.6        2/1/2:1                0d 00:10:16    2s        3
-------------------------------------------------------------------------------
Number of mrouters: 2
===============================================================================
*A:rbae_C#


*A:rbae_C# show service id 1 igmp-snooping mrouters detail
===============================================================================
IGMP Snooping Multicast Routers for service 1
===============================================================================
-------------------------------------------------------------------------------
MRouter 10.10.1.1
-------------------------------------------------------------------------------
Sap Id                  : 2/1/5:1
Expires                 : 17s
Up Time                 : 0d 00:00:32
Version                 : 3

General Query Interval  : 10s
Query Response Interval : 1.0s
```

```
Robust Count            : 2
-------------------------------------------------------------------------------
MRouter 10.20.1.6
-------------------------------------------------------------------------------
Sap Id                  : 2/1/2:1
Expires                 : 3s
Up Time                 : 0d 00:10:22
Version                 : 3

General Query Interval  : 2s
Query Response Interval : 1.0s
Robust Count            : 2
-------------------------------------------------------------------------------
Number of mrouters: 2
===============================================================================
*A:rbae_C#
```

mvr

Syntax **mvr**

Context show>service>id>igmp-snooping

Description Displays Multicast VPLS Registration (MVR) information.

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

Label	Description
MVR Admin State	Administrative state.
MVR Policy	Policy name.
Svc ID	The service identifier.
Sap/Sdp Id	Displays the SAP and SDP IDs of the service ID.
Oper State	Displays the operational state of the SAP and SDP IDs of the service ID.
Mrtr Port	Specifies if the port is a multicast router port.
From VPLS	Specifies from which VPLS the multicast streams corresponding to the groups learned via this SAP will be copied. If local, it is from its own VPLS.
Num Groups	Specifies the number of groups learned via this local SAP.

Sample Output

```
*A:ALA-1>show>service>id>snooping# mvr
===============================================================================
IGMP Snooping Multicast VPLS Registration info for service 10
```

```
===============================================================================
IGMP Snooping Admin State : Up

MVR Admin State           : Up
MVR Policy                : mvr-policy
-------------------------------------------------------------------------------
Local SAPs/SDPs
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp               Oper      From       Num Local
           Id                    State     VPLS       Groups
-------------------------------------------------------------------------------
100         sap:1/1/100:10       Up        Local      100
100         sap:1/1/100:20       Up        Local      100
-------------------------------------------------------------------------------
MVR SAPs (from-vpls=10)
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp               Oper      From       Num MVR
           Id                    State     VPLS       Groups
-------------------------------------------------------------------------------
20         sap:1/1/4:100         Up        10         100
30         sap:1/1/30:10.10      Up        10         100
===============================================================================
*A:ALA-1>show>service>id>snooping#
```

port-db

Syntax **port-db sap** *sap-id* [**detail**]
port-db sap *sap-id* **group** *grp-address*
port-db sdp *sdp-id*:*vc-id* [**detail**]
port-db sdp *sdp-id*:*vc-id* **group** *grp-address*

Context show>service>id>igmp-snooping

Description Displays information on the IGMP snooping port database for the VPLS service.

Parameters **group** *grp-ip-address* — Displays the IGMP snooping port database for a specific multicast group address.

sap *sap-id* — Displays the IGMP snooping port database for a specific SAP. The *sap-id* can be in one of the following formats:

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type*-*slot/mda*.*bundle-num*		
	bundle	keyword
	type	ima, ppp

	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp *sdp-id* **—** Displays only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Default For mesh SDPs only, all VC IDs.

Values 1 — 4294967295

group *grp-address* **—** Displays IGMP snooping statistics matching the specified group address.

source *ip-address* **—** Displays IGMP snooping statistics matching one particular source within the multicast group.

Sample Output

```
*A:ALA-1>show>service>id>snooping# port-db sap 1/1/2
===============================================================================
IGMP Snooping SAP 1/1/2 Port-DB for service 10
===============================================================================
Group Address    Mode     Type     Up Time        Expires   Num Sources
-------------------------------------------------------------------------------
225.0.0.1        include  dynamic  0d 00:04:44    0s        2
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================

*A:ALA-1>show>service>id>snooping# port-db sap 1/1/2 detail
===============================================================================
```

```
IGMP Snooping SAP 1/1/2 Port-DB for service 10
===============================================================================
IGMP Group 225.0.0.1
-------------------------------------------------------------------------------
Mode             : include              Type             : dynamic
Up Time          : 0d 00:04:57          Expires          : 0s
Compat Mode      : IGMP Version 3
V1 Host Expires  : 0s                   V2 Host Expires  : 0s
-------------------------------------------------------
Source Address  Up Time      Expires  Type     Fwd/Blk
-------------------------------------------------------
1.1.1.1         0d 00:04:57  20s      dynamic  Fwd
1.1.1.2         0d 00:04:57  20s      dynamic  Fwd
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================

*A:ALA-1>show>service>id>snooping#
```

proxy-db

Syntax **proxy-db** [**detail**]
proxy-db group *grp-address*

Context show>service>id>igmp-snooping

Description Displays information on the IGMP snooping proxy reporting database for the VPLS service.

Parameters **group grp-ip-address** — Displays the IGMP snooping proxy reporting database for a specific multicast group address.

```
*A:ALA-1>show>service>id>snooping# proxy-db
===============================================================================
IGMP Snooping Proxy-reporting DB for service 10
===============================================================================
Group Address    Mode     Up Time        Num Sources
-------------------------------------------------------------------------------
225.0.0.1        include  0d 00:05:40    2
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================

*A:ALA-1>show>service>id>snooping# proxy-db detail
===============================================================================
IGMP Snooping Proxy-reporting DB for service 10
===============================================================================
-------------------------------------------------------------------------------
IGMP Group 225.0.0.1
-------------------------------------------------------------------------------
Up Time : 0d 00:05:54                   Mode : include
-----------------------------
Source Address  Up Time
-----------------------------
1.1.1.1         0d 00:05:54
1.1.1.2         0d 00:05:54
-------------------------------------------------------------------------------
Number of groups: 1
```

```
===============================================================================

*A:ALA-1>show>service>id>snooping#
```

querier

Syntax **querier**

Context show>service>id>igmp-snooping

Description Displays information on the IGMP snooping queriers for the VPLS service.

```
*A:ALA-1>show>service>id>snooping# querier
===============================================================================
IGMP Snooping Querier info for service 10
===============================================================================
Sap Id                  : 1/1/1
IP Address              : 10.10.10.1
Expires                 : 6s
Up Time                 : 0d 00:56:50
Version                 : 3

General Query Interval  : 5s
Query Response Interval : 2.0s
Robust Count            : 2
===============================================================================
*A:ALA-1>show>service>id>snooping#
```

static

Syntax **static** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>igmp-snooping

Description Displays information on static IGMP snooping source groups for the VPLS service.

Parameters **sap** *sap-id* — Displays static IGMP snooping source groups for a specific SAP. The *sap-id* can be in one of the following formats:

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port*.*channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64

bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp *sdp-id* **—** Displays the IGMP snooping source groups for a specific spoke or mesh SDP.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Default For mesh SDPs only, all VC IDs.

Values 1 — 4294967295

```
*A:ALA-1>show>service>id>snooping# static
===============================================================================
-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups for SAP 1/1/2
-------------------------------------------------------------------------------
Source           Group
-------------------------------------------------------------------------------
*                225.0.0.2
*                225.0.0.3
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 2

-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups for SDP 10:10
-------------------------------------------------------------------------------
Source           Group
-------------------------------------------------------------------------------
1.1.1.1          225.0.0.10
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 1
```

```
===============================================================================

*A:ALA-1>show>service>id>snooping#
```

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>igmp-snooping

Description Displays IGMP snooping statistics for the VPLS service.

Parameters **sap** *sap-id* **—** Displays IGMP snooping statistics for a specific SAP.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp *sdp-id* **—** Displays the IGMP snooping statistics for a specific spoke or mesh SDP.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Default For mesh SDPs only, all VC IDs.

Values 1 — 4294967295

```
*A:ALA-1>show>service>id>snooping# statistics
===============================================================================
IGMP Snooping Statistics for service 1
===============================================================================
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         4             0             4
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              0             0             0
V3 Reports              0             0             0
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length              : 0
Bad IP Checksum         : 0
Bad IGMP Checksum       : 0
Bad Encoding            : 0
No Router Alert         : 0
Zero Source IP          : 0

Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0

MVR From VPLS Cfg Drops : 0
MVR To SAP Cfg Drops    : 0
===============================================================================
*A:ALA-1>show>service>id>snooping#
```

egress-replication

Syntax **egress-replication**

Context show

Description This command enables the context to display egress flooding information for a VPLS service context on a given MDA. A VPLS service context supports both Layer 2 and Layer 3 flooding modes. The Layer 2 flooding mode is used for broadcast, Layer 2 multicast and unknown destination MAC addressed packets. All available interfaces (SAP, spoke-SDP and mesh-SDP) that reside on an egress forwarding complex are included in the egress list except for SAPs that are defined in a residential split horizon group (Layer 2 flooding is not permitted on residential SAPs). The Layer 3 flooding

mode is used for VPLS interfaces participating in IGMP snooping and is represented by an IP multicast [s,g] record.

vpls

Syntax **vpls** *vpls-service-id* **mda** *card/slot*
vpls *vpls-service-id* **mda** *card/slot* [**igmp-record** *group ip-address* {**source** *ip-address* | **starg**}]

Context show>egress-replication

Description The **vpls** *vpls-service-id* **mda** *slot/mda* command displays the flooding list used by the Layer 2 flooding mode for the VPLS service on the specified MDA. The Layer 2 flooding list is limited to SAPs, spoke-SDP and mesh-SDP bindings that exist on the egress forwarding complex serviced by the specified MDA. For the 10G IOM, two MDAs share the same egress forwarding plane. In this case the Layer 2 flooding list will contain destinations for both MDAs (if entries exist). The only VPLS interfaces that will not be included in the list are residential SAPs because Layer 2 replication is not permitted to a residential SAP. A packet processed by the egress Layer 2 flooding list may not be replicated to each destination. A packet will not be replicated to an interface on the Layer 2 flooding list because of the following:

The ingress interface is the same as egress interface (source squelching rule)

- The ingress interface split horizon group is the same as the egress interface (residential bridging rule)
- The egress interface is down or blocking
- The packet matches a discard event while processing that destination interface
- An egress MTU violation occurs for the destination interface

Destination SAPs in the list may be displayed in a chain context representing common replication behavior. All SAPs in a single chain are processed a single time through the egress forwarding plane. If a discard decision is made for the first SAP in the chain, no replication processing is done for any of the chain members. If the forwarding plane decides to replicate the first SAP in the chain, it will replicate to all SAPs in the chain.

The **vpls** *vpls-service-id* **mda card/slot igmp-record** *grp-address* {**source** *source-ip-address* | **starg**} command displays the IGMP record based flooding list for the *vpls-service-id* on the specified MDA. Unlike the Layer 2 flooding list for the VPLS context, an IGMP record list may contain interfaces from other VPLS contexts due to MVR (Multicast VPLS Registration) events on the individual VPLS interfaces. VPLS interfaces in other VPLS contexts become associated with the specified vpls-service-id based on the MVR from-vpls definition. Another difference between the VPLS Layer 2 flooding list and IGMP lists is that many IGMP lists may exist (each associated with a different [s,g] record) and the lists may contain residential SAPs. The SAP chaining and replication behavior is similar to the VPLS Layer 2 flooding list.

IP multicast packets ingressing the vpls-service-id must match either a [*,g] or [s,g] record to be associated with the record's egress IP multicast IGMP flooding list. A [*,g] record will match any ingress IP multicast packet destined to the class D destination IP address represented by "g". An [s,g] record will match any ingress IP multicast packet with a source IP address matching "s" and a destination IP address matching "g". In the case that a packet could match both a [*,g] and [s,g] record, the [s,g] record takes precedence. Each [*,g] and [s,g] record has its own IGMP flooding list. The list will only appear on an egress forwarding plane (MDA) when a member of the list (VPLS interface) exists on the forwarding plane.

IGMP Commands

group

Syntax **group** [*grp-ip-address*]

Context show>router>igmp

Description This command displays the multicast group and (s, g) addresses. If no *grp-ip-address* parameters are specified then all IGMP group, (*, g) and (s, g) addresses are displayed.

Parameters *grp-ip-address* — Displays specific multicast group addresses.

Output **IGMP Group Output —** The following table describes the output fields for IGMP group information.

Label	Description
IGMP Groups	Displays the IP multicast sources corresponding to the IP multicast groups which are statically configured.
Fwd List	Displays the list of interfaces in the forward list.
Blk List	Displays the list of interfaces in the block list.

Sample Output

```
A:NYC# show router igmp group
===============================================================================
IGMP Groups
===============================================================================
(*,224.24.24.24)                        Up Time : 0d 05:21:38
    Fwd List  : nyc-vlc

(*,239.255.255.250)                     Up Time : 0d 05:21:38
    Fwd List  : nyc-vlc
-------------------------------------------------------------------------------
(*,G)/(S,G) Entries : 2
===============================================================================
A:NYC#


A:NYC# show router igmp group 224.24.24.24
===============================================================================
IGMP Groups
===============================================================================
(*,224.24.24.24)                        Up Time : 0d 05:23:23
    Fwd List  : nyc-vlc
-------------------------------------------------------------------------------
(*,G)/(S,G) Entries : 1
===============================================================================
A:NYC#
```

ssm-translate

Syntax **ssm-translate**

Context show>router>igmp

Description This command displays IGMP SSM translate configuration information.

Output **GMP Interface Output —** The following table provides IGMP field descriptions

Label	Description
Group Range	Displays the address ranges of the multicast groups for which this router can be an RP.
Source	Displays the unicast address that sends data on an interface.
SSM Translate Entries	Displays the total number of SSM translate entries.

```
A:ALA-48>config>router>igmp# show router igmp ssm-translate
=======================================================
IGMP SSM Translate Entries
=======================================================
Group Range                              Source
-------------------------------------------------------
<224.0.1.0 - 224.0.1.255>                1.1.1.1
<225.1.0.0 - 225.240.3.57>               2.2.2.2
<239.255.255.0 - 239.255.255.255>        3.3.3.3
-------------------------------------------------------
SSM Translate Entries : 3
=======================================================
A:ALA-48>config>router>igmp#
```

interface

Syntax **interface** [*ip-int-name* | *ip-address*] [**group**] [*grp-address*] [**detail**]

Context show>router>igmp

Description This command displays IGMP interface information.

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.

group *grp-address* — Only displays IP multicast group address for which this entry contains information.

detail — Displays detailed IP interface information along with the source group information learned on that interface.

Output **IGMP Interface Output —** The following table provides IGMP field descriptions

Label	Description
Interface	Specifies the interfaces that participates in the IGMP protocol.
Adm Admin Status	Displays the administrative state for the IGMP protocol on this interface.
Oper Oper Status	Displays the current operational state of IGMP protocol on the interface.
Querier	Displays the address of the IGMP querier on the IP subnet to which the interface is attached.
Querier Up Time	Displays the time since the querier was last elected as querier.
Querier Expiry Timer	Displays the time remaining before the querier ages out. If the querier is the local interface address, the value will be zero.
Cfg/Opr Version Admin/Oper version	Cfg – The configured version of IGMP running on this interface. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN. Opr – The operational version of IGMP running on this interface. If the cfg value is 3 but all of the routers in the local subnet of this interface use IGMP version v1 or v2, the operational version will be v1 or v2.
Num Groups	The number of multicast groups which have been learned by the router on the interface.
Policy	Specifies the policy that is to be applied on the interface.
Group Address	Specifies the IP multicast group address for which this entry contains information.
Up Time	Specifies the time since this source group entry got created.
Last Reporter	Specifies the IP address of the source of the last membership report received for this IP Multicast group address on this interface. If no membership report has been received, this object has the value 0.0.0.0.
Mode	The mode is based on the type of membership report(s) received on the interface for the group. In the 'include' mode, reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the source-list parameter of the IGMP membership report. In 'exclude' mode, reception of packets sent to the given multicast address is requested from all IP source addresses except those listed in the source-list parameter.

Label	Description (Continued)
V1 Host Timer	The time remaining until the local router will assume that there are no longer any IGMP version 1 members on the IP subnet attached to this interface. Upon hearing any IGMPv1 Membership Report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv2 Leave messages for this group that it receives on this interface.
V2 Host Timer	The time remaining until the local router will assume that there are no longer any IGMP version 2 members on the IP subnet attached to this interface. Upon hearing any IGMPv2 Membership Report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv3 Leave messages for this group that it receives on this interface.
Type	Indicates how this group entry was learned. If this group entry was learned by IGMP, it will be set to 'dynamic'. For statically configured groups, the value will be set to 'static'.
Compat Mode	Used in order for routers to be compatible with older version routers. IGMPv3 hosts MUST operate in version 1 and version 2 compatibility modes. IGMPv3 hosts MUST keep state per local interface regarding the compatibility mode of each attached network. A host's compatibility mode is determined from the Host Compatibility Mode variable which can be in one of three states: IGMPv1, IGMPv2 or IGMPv3. This variable is kept per interface and is dependent on the version of General Queries heard on that interface as well as the Older Version Querier Present timers for the interface.

Sample Output

```
A:BA# show router igmp interface
===============================================================================
IGMP Interfaces
===============================================================================
Interface               Adm  Oper Querier         Cfg/Opr Num     Policy
                                                  Version Groups
-------------------------------------------------------------------------------
IGMP_to_CE              Up   Up   11.1.1.1        1/1     3       igmppol
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
A:BA#


A:BA# show router 100 igmp interface IGMP_to_CE
===============================================================================
IGMP Interface IGMP_to_CE
===============================================================================
Interface               Adm  Oper Querier         Cfg/Opr Num     Policy
                                                  Version Groups
-------------------------------------------------------------------------------
IGMP_to_CE              Up   Up   11.1.1.1        1/1     3       igmppol
-------------------------------------------------------------------------------
```

```
Interfaces : 1
===============================================================================
A:BA#


A:BA# show router 100 igmp interface 11.1.1.1
===============================================================================
IGMP Interface 11.1.1.1
===============================================================================
Interface               Adm  Oper Querier         Cfg/Opr Num     Policy
                                                  Version Groups
-------------------------------------------------------------------------------
IGMP_to_CE              Up   Up   11.1.1.1        1/1     3       igmppol
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
A:BA#


A:BA# show router 100 igmp interface IGMP_to_CE group 227.1.1.1
===============================================================================
IGMP Interface IGMP_to_CE
===============================================================================
Interface               Adm  Oper Querier         Cfg/Opr Num     Policy
                                                  Version Groups
-------------------------------------------------------------------------------
IGMP_to_CE              Up   Up   11.1.1.1        1/1     3       igmppol
-------------------------------------------------------------------------------
IGMP Group
-------------------------------------------------------------------------------
Group Address : 227.1.1.1              Up Time       : 0d 00:03:52
Interface     : IGMP_to_CE             Expires       : never
Last Reporter : 0.0.0.0                Mode          : exclude
V1 Host Timer : Not running            Type          : static
V2 Host Timer : Not running            Compat Mode   : IGMP Version 3
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================


A:BA# show router 100 igmp interface IGMP_to_CE group 227.1.1.1 detail
===============================================================================
IGMP Interface IGMP_to_CE
===============================================================================
Interface          : IGMP_to_CE
Admin Status       : Up                Oper Status        : Up
Querier            : 11.1.1.1          Querier Up Time    : 0d 00:04:01
Querier Expiry Time: N/A               Time for next query: 0d 00:13:42
Admin/Oper version : 1/1               Num Groups         : 3
Policy             : igmppol           Subnet Check       : Disabled
Max Groups Allowed : 16000             Max Groups Till Now: 3
MCAC Policy Name   :                   MCAC Const Adm St  : Enable
MCAC Max Unconst BW: no limit          MCAC Max Mand BW   : no limit
MCAC In use Mand BW: 0                 MCAC Avail Mand BW : unlimited
MCAC In use Opnl BW: 0                 MCAC Avail Opnl BW : unlimited
-------------------------------------------------------------------------------
IGMP Group
-------------------------------------------------------------------------------
Group Address : 227.1.1.1              Up Time       : 0d 00:04:02
Interface     : IGMP_to_CE             Expires       : never
Last Reporter : 0.0.0.0                Mode          : exclude
V1 Host Timer : Not running            Type          : static
V2 Host Timer : Not running            Compat Mode   : IGMP Version 3
```

```
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
A:BA#
```

static

Syntax **static** [*ip-int-name* | *ip-addr*]

Context show>router>igmp

Description This command displays static IGMP, (*, g) (s, g) information.

Parameters *ip-int-name* — Only displays the information associated with the specified IP interface name.

ip-addr — Only displays the information associated with the specified IP address.

Output **Static IGMP Output —** The following table provides static IGMP field descriptions

Label	Description
Source	Displays entries which represents a source address from which receivers are interested/not interested in receiving multicast traffic.
Group	Displays the IP multicast group address for which this entry contains information.
Interface	Displays the interface name.

Sample Output

```
A:BA# show router 100 igmp static
===============================================================
IGMP Static Group Source
===============================================================
Source          Group           Interface
---------------------------------------------------------------
11.11.11.11     226.136.22.3    IGMP_to_CE
*               227.1.1.1       IGMP_to_CE
22.22.22.22     239.255.255.255 IGMP_to_CE
---------------------------------------------------------------
Static (*,G)/(S,G) Entries : 3
===============================================================
A:BA#
```

statistics

Syntax **statistics** [*ip-int-name* | *ip-addr*ess]

Context show>router>igmp

Description This command displays IGMP statistics information.

Parameters *ip-int-name* — Only displays the information associated with the specified IP interface name.

ip-addr — Only displays the information associated with the specified IP address.

Output **IGMP Statistics Output —** The following table provides statistical IGMP field descriptions

Label	Description
IGMP Interface Statistics	The section listing the IGMP statistics for a particular interface.
Message Type	Queries — The number of IGMP general queries transmitted or received on this interface.
	Report — The total number of IGMP V1, V2, or V3 reports transmitted or received on this interface.
	Leaves — The total number of IGMP leaves transmitted on this interface.
Received	Column that displays the total number of IGMP packets received on this interface.
Transmitted	Column that displays the total number of IGMP packets transmitted from this interface.
General Interface Statistics	The section listing the general IGMP statistics.
Bad Length	Displays the total number of IGMP packets with bad length received on this interface.
Bad Checksum	Displays the total number of IGMP packets with bad checksum received on this interface.
Unknown Type	Displays the total number of IGMP packets with unknown type received on this interface.
Bad Receive If	Displays the total number of IGMP packets incorrectly received on this interface.
Rx Non Local	Displays the total number of IGMP packets received from a non-local sender.
Rx Wrong Version	Displays the total number of IGMP packets with wrong versions received on this interface.
Policy Drops	Displays the total number of times IGMP protocol instance matched the host IP address or group/source addresses specified in the import policy.
No Router Alert	Displays the total number of IGMPv3 packets received on this interface which did not have the router alert flag set.

Sample Output

```
A:BA# show router 100 igmp statistics
```

```
===============================================
IGMP Interface Statistics
===============================================
Message Type        Received       Transmitted
-----------------------------------------------
Queries             0              5
Report V1           0              0
Report V2           0              0
Report V3           0              0
Leaves              0              0
-----------------------------------------------
General Interface Statistics
-----------------------------------------------
Bad Length        : 0
Bad Checksum      : 0
Unknown Type      : 0
Bad Receive If    : 0
Rx Non Local      : 0
Rx Wrong Version  : 0
Policy Drops      : 0
No Router Alert   : 0
Rx Bad Encodings  : 0
Rx Pkt Drops      : 0
-----------------------------------------------
Source Group Statistics
-----------------------------------------------
(S,G)             : 2
(*,G)             : 1
===============================================
A:BA#
```

status

Syntax **status**

Context show>router>igmp

Description This command displays IGMP status information.
If IGMP is not enabled, the following message appears:

```
A:NYC# show router igmp status
MINOR: CLI IGMP is not configured.
A:NYC#
```

Output **IGMP Status Output —** The following table provides IGMP status field descriptions

Label	Description
Admin State	Displays the administrative status of IGMP.
Oper State	Displays the current operating state of this IGMP protocol instance on this router.
Query Interval	The frequency at which IGMP query packets are transmitted.

Label	Description (Continued)
Last Member Query Interval	The maximum response time inserted into Group-Specific Queries sent in response to Leave Group messages, and is also the amount of time between Group-Specific Query messages.
Query Response Interval	The maximum query response time advertised in IGMPv2 queries.
Robust Count	Displays the number of times the router will retry a query.

Sample Output

```
A:BA# show router 100 igmp status
===============================================================================
IGMP Status
===============================================================================
Admin State                       : Up
Oper State                        : Up
Query Interval                    : 1024
Last Member Query Interval        : 1024
Query Response Interval           : 1023
Robust Count                      : 10
===============================================================================
A:BA
```

Clear Commands

id

Syntax **id** *service-id*

Context clear>service
clear>service>statistics

Description This command clears commands for a specific service.

Parameters *service-id* — The ID that uniquely identifies a service.

authentication

Syntax **authentication**

Context clear>service>id

Description Enters the context to clear session authentication information.

statistics

Syntax **statistics**

Context clear>service>id>authentication

Description Clears session authentication statistics for this service.

fdb

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

Context clear>service>id

Description Clears FDB entries for the service.

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

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

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

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.

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

Values	sdp-id[:vc-id]	*sdp-id*	1 — 17407
		vc-id	1 — 4294967295
	sdp-id:vc-id	sdp-id	1 — 17407
		vc-id	1 — 4294967295

mesh-sdp

Syntax **mesh-sdp** *sdp-id*[**:***vc-id*] **ingress-vc-label**

Context clear>service>id

Description Clears and resets the mesh SDP bindings for the service.

Parameters *sdp-id* — The mesh SDP ID to be reset.

Values 1 — 17407

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

Default All VC IDs on the SDP ID.

Values 1 — 4294967295

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

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

Values 1 — 4294967295

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.

all — Clears all SAP queue statistics and STP statistics.

counters — Clears all queue statistics associated with the SAP.

stp — Clears all STP statistics associated with the SAP.

sdp

Syntax **sdp** *sdp-id* [**keep-alive**]

Context clear>service>statistics

Description Clears keepalive statistics associated with the SDP ID.

Parameters *sdp-id* — The SDP ID for which to clear statistics.

Values 1 — 17407

keep-alive — Clears the keep-alive history associated with this SDP ID.

counters

Syntax **counters**

Context clear>service>statistics>id

Description Clears all traffic queue counters associated with the service ID.

sap

Syntax **sap** *sap-id* {**all** | **counters** | **stp**}

Context clear>service>statistics>id

Description Clears statistics for the SAP bound to the service.

Parameters *sap-id* — The SAP ID for which to clear statistics.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword

	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type*-*bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[:*vc-id*] {**all** | **counters** | **stp**}

Context clear>service>statistics>id

Description 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.

detected-protocols

Syntax **detected-protocols** { **all** | **sap** *sap-id* | **spoke-sdp** *sdp-id*[*:vc-id*]}

Context clear>service>id>stp

Description RSTP automatically falls back to STP mode when it receives an STP BPDU. The **clear detected-protocols** command forces the system to revert to the default RSTP mode on the SAP or spoke SDP.

Parameters **all —** Clears all detected protocol statistics.

sap-id — Clears the specified lease state SAP information.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255

vci	1, 2, 5 — 65535
dlci	16 — 1022

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

lease-state

Syntax **lease-state** [**no-dhcp-release**]
lease-state ip-address *ip-address* [**no-dhcp-release**]
lease-state mac *ieee-address* **no-dhcp-release**
lease-state sap *sap-id* [**no-dhcp-release**]
lease-state sdp *sdp-id:vc-id* [**no-dhcp-release**]

Context clear>service>id>dhcp

Description This command clears DHCP lease state information.

Parameters **no-dhcp-release —** Specifies that the node will clear the state without sending the DHCP release message.

ip-address *ip-address* **—** Clears the DHCP IP address lease state information. 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 /30 subnets).

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

sap *sap-id* **—** Clears DHCP SAP lease state information.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword

	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp *sdp-id:vc-id* **—** *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

statistics

Syntax **statistics** [**sap** *sap-id* **| sdp** [*sdp-id*[*:vc-id*] | **interface** [*ip-address* | *ip-int-name*]]

Context clear>service>id>dhcp

Description Clears DHCP statistics for this service.

Parameters *sap-id* — Clears the specified SAP statistics.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*

port-id	*slot/mda/port*[*.channel*]	
aps-id	aps-*group-id*[*.channel*]	
	aps	keyword
	group-id	1 — 64

bundle-*type-slot/mda.bundle-num*

	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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

interface *ip-int-name* **—** Clears the statistics for the IP interface with the specified name.

interface *ip-addr* **—** Clears the statistics for the IP interface with the specified IP address.

port-db

Syntax **port-db** {**sap** *sap-id* | **sdp** *sdp-id:vc-id*} [**group** *grp-address* [**source** *ip-address*]]

Context clear>service>id>igmp-snooping

Description Clears the information on the IGMP snooping port database for the VPLS service.

Parameters **sap** *sap-id* **—** Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200

qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp *sdp-id* **—** Clears only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 — 17407

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

Default For mesh SDPs only, all VC IDs.

Values 1 — 4294967295

group *grp-address* **—** Clears IGMP snooping statistics matching the specified group address.

source *ip-address* **—** Clears IGMP snooping statistics matching one particular source within the multicast group.

querier

Syntax **querier**

Context clear>service>id>igmp-snooping

Description Clears the information on the IGMP snooping queriers for the VPLS service.

statistics

Syntax **statistics** {**all** | **sap** *sap-id* | **sdp** *sdp-id:vc-id*}]

Context clear>service>id>igmp-snooping

Description Clears IGMP snooping statistics for the VPLS service.

Parameters **sap** *sap-id* — Clears the IGMP snooping information on the specified SAP.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp *sdp-id* — Clears only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 — 17407

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

Default For mesh SDPs only, all VC IDs.

Values 1 — 4294967295

mfib

Syntax **mfib**

Context clear>service>id>

Description Enter the context to clear multicast FIB info for the VPLS service.

statistics

Syntax **statistics** {**all** | **group** *grp-address*}

Context clear>service>id>mfib

Description Clears multicast FIB statistics for the VPLS service.

Parameters *grp-address* — Specifies an IGMP multicast group address that receives data on an interface.

arp

Syntax **arp** {**all** | *ip-address*}
arp interface [*ip-int-name* | *ip-address*]

Context clear>router

Description This command clears all or specific ARP entries.

The scope of ARP cache entries cleared depends on the command line option(s) specified.

Parameters **all —** Clears all ARP cache entries.

ip-address — Clears the ARP cache entry for the specified IP address.

interface *ip-int-name* **—** Clears all ARP cache entries for the IP interface with the specified name.

interface *ip-addr* **—** Clears all ARP cache entries for the specified IP interface with the specified IP address.

dhcp

Syntax **dhcp**

Context clear>router

Description This command enables the context to clear and reset DHCP entities.

statistics

Syntax **statistics** [**interface** *ip-int-name* | *ip-address*]

Context clear>router>dhcp

Description Clears DHCP statistics.

interface *ip-int-name* **—** Clears the statistics for the IP interface with the specified name.

interface *ip-addr* **—** Clears the statistics for the IP interface with the specified IP address.

Debug Commands

id

Syntax **id** *service-id*

Context debug>service

Description This command debugs commands for a specific service.

Parameters *service-id* — The ID that uniquely identifies a service.

stp

Syntax **stp**

Context debug>service

Description This command enables the context for debugging STP.

all-events

Syntax **all-events**

Context debug>service>stp

Description This command enables STP debugging for all events.

bpdu

Syntax [**no**] **bpdu**

Context debug>service>stp

Description This command enables STP debugging for received and transmitted BPDUs.

core-connectivity

Syntax [**no**] **core-connectivity**

Context debug>service>stp

Description This command enables STP debugging for core connectivity.

exception

Syntax [**no**] **exception**

Context debug>service>stp

Description This command enables STP debugging for exceptions.

fsm-state-changes

Syntax [**no**] **fsm-state-changes**

Context debug>service>stp

Description This command enables STP debugging for FSM state changes.

fsm-timers

Syntax [**no**] **fsm-timers**

Context debug>service>stp

Description This command enables STP debugging for FSM timer changes.

port-role

Syntax [**no**] **port-role**

Context debug>service>stp

Description This command enables STP debugging for changes in port roles.

port-state

Syntax [**no**] **port-state**

Context debug>service>stp

Description This command enables STP debugging for port states.

sap

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

Context debug>service>stp

Description This command enables STP debugging for a specific SAP.

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>stp

Description This command enables STP debugging for a specific SDP.

interface

Syntax [**no**] **interface** [*ip-int-name* | *ip-address*]

Context debug>router>igmp

Description This command enables debugging on the IGMP interface.

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.

mcs

Syntax [**no**] **mcs** [*ip-int-name]*

Context debug>router>igmp

Description This command enables debugging for IGMP MCS.

Parameters *ip-int-name* — Only displays the information associated with the specified IP interface name.

misc

Syntax [**no**] **misc**

Context debug>router>igmp

Description This command enables debugging for IGMP miscellany.

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.

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.

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

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 access point to the subscriber network. Multiple IES services are created to segregate subscriber-owned IP interfaces.

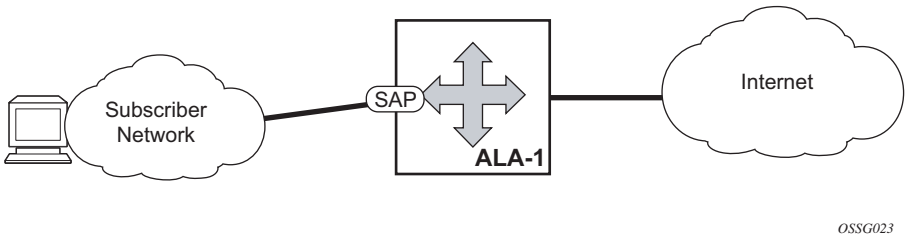

Figure 1: Internet Enhanced Service

The IES service provides Internet connectivity. Other features include:

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

Note: Refer to the 7750 SR OS Triple Play Guide for information about how subscriber group-interfaces function in the Routed Central Office model.

IES Features

This section describes various of the general 7750 SR 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:

- VRRP - for IES services with more than one IP interface
- Cflowd
- Secondary IP addresses
- ICMP Options

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

- Unnumbered interfaces
- NTP broadcast receipt

Subscriber Interfaces

Subscriber interfaces are composed of a combination of two key technologies, subscriber interfaces and group interfaces. While the subscriber interface defines the subscriber subnets, the group interfaces are responsible for aggregating the SAPs.

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

SAPs

Encapsulations

The following SAP encapsulations are supported on the 7750 SR IES service:

- Ethernet null
- Ethernet dot1q
- SONET/SDH IPCP
- SONET/SDH BCP-null
- SONET/SDH BCP-dot1q
- SONET/SDH ATM
- ATM - LLC SNAP or VC-MUX

ATM SAP Encapsulations for IES

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

The following ATM encapsulation and transport modes are supported:

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

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.

QoS Policies

When applied to 7750 SR 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 L2 or L3 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.

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 Service Distribution Points (SDPs) on page 41 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 50 depicts traffic terminating on a specific IES or VPRN service that is identified by the *sdp-id* and VC label present in the the service packet.

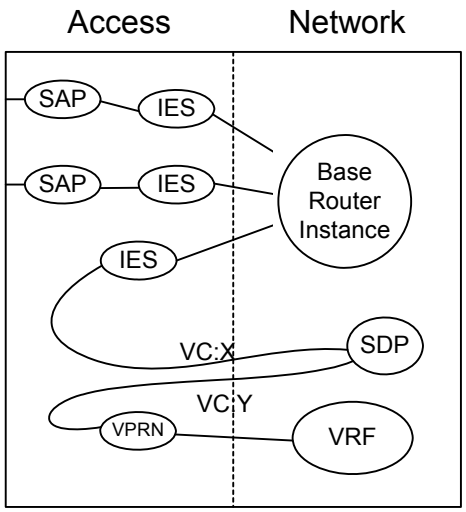

Figure 2: SDP-ID and VC Label Service Identifiers

Figure 3: IES Spoke-SDP Termination

Figure 51 depicts a spoke SDP terminating directly into an IES. In this case, a spoke SDP could be tied to an Epipe or HVPLS service. There is no configuration required on the PE connected to the CE.

Note that all the routing protocols, including multicast, that are supported by IES are supported for spoke-sdp termination.

SRRP

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

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

For proper operation, each subscriber subnet associated with the SRRP instance must have a gw-address defined. The SRRP instance cannot be activated (no shutdown) unless each subscriber subnet associated with the group IP interface has an SRRP gateway IP address. Once the SRRP instance is activated, new subscriber subnets cannot be added without a corresponding SRRP gateway IP address. Table 22 describes how the SRRP instance state is used to manage access to subscriber hosts associated with the group IP interface.

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

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

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

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

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

Table 22 lists the SRRP's state effect on subscriber hosts associated with group IP interfaces.

Table 1: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface

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

Table 1: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface (Continued)

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

Table 1: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface (Continued)

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

Table 1: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface (Continued)

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

SRRP Messaging

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

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

The vr-id field has been expanded to support an SRRP instance ID of 32 bits.

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

Unlike VRRP, only two nodes may participate in an SRRP instance due the explicit association between the SRRP instance group IP interface, the associated redundant IP interface and the multi-chassis synchronization (MCS) peering. Since only two nodes are participating, the VRRP skew

timer is not utilized when waiting to enter the master state. Also, SRRP always preempts when the local priority is better than the current master and the backup SRRP instance always inherits the master's advertisement interval from the SRRP advertisement messaging.

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

The SRRP advertisement message is always evaluated to see if it has higher priority than the SRRP advertisement that would be sent by the local node. If the advertised priority is equal to the current local priority, the source IP address of the received SRRP advertisement is used as a tie breaker. The node with the lowest IP address is considered to have the highest priority.

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

SRRP and Multi-Chassis Synchronization

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

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

- The SRRP instance MCS key
- Containing service type and ID
- Containing subscriber IP interface name
- Subscriber subnet information

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

SRRP Instance

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

- SRRP Instance MCS Key on page 676
- Containing Service Type and ID on page 677
- Containing Subscriber IP Interface Name on page 677
- Subscriber Subnet Information on page 677

SRRP Instance MCS Key

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

Containing Service Type and ID

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

Containing Subscriber IP Interface Name

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

Subscriber Subnet Information

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

Containing Group IP Interface Information

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

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

Remote Redundant IP Interface Mismatch

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

Remote Sending Redundant IP Interface Unavailable

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

Remote SRRP Advertisement SAP Non-existent

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

Remote Sending Local Receive SRRP Advertisement SAP Unavailable

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

Local and Remote Dual Master Detected

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

Subscriber Subnet Owned IP Address Connectivity

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

Subscriber Subnet SRRP Gateway IP Address Connectivity

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

Receive SRRP Advertisement SAP and Anti-Spoof

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

Configuring an IES Service with CLI

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

Topics in this section include:

List of Commands

Table 23 lists all the service configuration commands indicating the configuration level at which each command is implemented with a short command description. The command list is organized in the following task-oriented manner:

- Create an IP address range reserved for IES
- Configure IES service parameters
  - → Configure IES subscriber interface parameters on page 682
  - → Configure IES subscriber interface group interface parameters on page 683
  - → Configure interface parameters
    - – Configure IES interface ICMP parameters
    - – Configure IES interface IPv6 parameters
    - – Configure IES interface spoke-SDP parameters
    - – Configure VRRP parameters
- Configure SAP parameters
  - → Configure SAP ATM parameters

Table 1: CLI Commands to Configure IES Service Parameters

Command	Description	Page
Create an IP address range reserved for IES		
`config>router`		
`service-prefix`	Creates an IP address range reserved for IES to provide a mechanism to reserve one or more address ranges for services.	
Configure IES service parameters		691
`config>service`		
`ies`	Creates an IES service instance and specifies the customer ID number to be associated with the service.	719
`customer`	Specifies the customer ID number to be associated with the service.	720
`no shutdown`	Administratively enables the service.	717
Configure IES subscriber interface parameters		
`config>service>ies>sub-if`		
`address`	Assigns an IP address, IP subnet, and broadcast address format to an IES IP router interface.	729
`dhcp`	Configures DHCP parameters for the subscriber interface.	742
`group-interface`	Enables the context to configure a group interface.	

Table 1: CLI Commands to Configure IES Service Parameters (Continued)

Command	Description	Page
Configure IES subscriber interface group interface parameters		
`config>service>ies>sub-if>grp-if`		
`arp-populate`	Disables dynamic learning of ARP entries.	739
`arp-timeout`	Configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table.	733
`authentication-policy`	Assigns a RADIUS authentication policy to the interface.	722
`description`	Creates a text description stored in the configuration file for a configuration context.	718
`dhcp`	Configures DHCP parameters for this interface	742
`host-connectivity-verify`	Enables host connectivity verification for all hosts on this interface.	769
`local-proxy-arp`	Enables local proxy ARP on the interface.	734
`mac`	Assigns a specific MAC address to an IES IP interface.	735
`proxy-arp-policy`	Configures a proxy ARP policy for the interface.	738
`remote-proxy-arp`	Enables remote proxy ARP on the interface.	739
`tos-marking-state`	Changes the default trusted state to a non-trusted state.	737
Configure IES group interface SRRP parameters		
`config>service>ies>sub-if>grp-if>srrp`		
`gw-mac`	Overrides the default SRRP gateway MAC address used by the SRRP instance.	723
`keep-alive-interval`	defines the interval between SRRP advertisement messages sent when operating in the master state.	724
`message-path`	defines a specific SAP for SRRP in-band messaging.	724
`policy`	associates one or more VRRP policies with the SRRP instance.	726
`priority`	overrides the default base priority for the SRRP instance.	726
`no shutdown`	Administratively enables the entity.	717
Configure interface parameters		693
`config>service>ies>interface` *`interface-name`*		
`active-cpm-protocols`	Enables or disables CPM protocols on this interface.	729
`address`	Assigns an IP address, IP subnet and broadcast address format to an IES or IES IP router interface.	729
`allow-directed-broadcasts`	Enables the forwarding of directed broadcasts out of the IP interface.	731
`arp-populate`	Enables populating static and dynamic hosts into the system ARP cache.	739

Table 1: CLI Commands to Configure IES Service Parameters (Continued)

Command	Description	Page
arp-timeout	Configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table.	733
authentication-policy	Assigns a RADIUS authentication policy to the interface	722
bfd	Specifies the BFD parameters for the associated IP interface	733
cflowd	Configures *cflowd* collection and analysis on the (routeable) interface.	734
host	Creates a static host for the SAP. Applications within the system that make use of static host entries include anti-spoof and source MAC population into the VPLS forwarding database.	740
host-connectivity-verify	Enables or disable host connectivity verification for all hosts on this interface.	769
ip-mtu	Configures the interface IP maximum transmit unit.	735
loopback	Specifies that the associated interface is a loopback interface that has no associated physical interface.	735
mac	Assigns a specific MAC address to an IES or IES IP interface.	735
proxy-arp	Enables a proxy ARP policy for the interface.	738
secondary	Assigns an secondary IP address/IP subnet/broadcast address format to the interface.	736
static-arp	Configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance.	737
tos-marking-state	Configures the ToS field to a trusted state or a non-trusted state.	737
no shutdown	Administratively enables the service.	717
Configure IES DHCP parameters		
config>service>ies>if>dhcp		
config>service>ies>sub-if>grp-if>dhcp		
dhcp	Configures DHCP parameters for this interface.	742
description	Specifies a text string description for DHCP on this interface.	718
gi-address	Configures the gateway interface address for the DHCP relay.	747
lease-populate	Enables dynamic host lease state management for IP interfaces.	748
option	Configures DHCP Option 82 processing for this interface.	743
server	Configures the DHCP server IP address.	747
no shutdown	Enables DHCP on this interface.	717
trusted	Enables relaying of untrusted packets.	747
config>service>ies>if>dhcp>option		
action	Configures the DHCP relay reforwarding policy action.	742

Table 1: CLI Commands to Configure IES Service Parameters (Continued)

Command	Description	Page
`circuit-id`	Enables sending of the interface index in the circuit-id suboption of the DHCP relay packet.	742
`remote-id`	Enables sending of the remote MAC address in the remote-id suboption of the DHCP relay packet.	744
`match-circuit-id`	Enables Option 82 circuit ID on relayed DHCP packet matching only in the **config>service>ies>sub-if>grp-if>dhcp** context.	743
Configure IES interface ICMP parameters		
`config>service>ies>interface`		
`config>service>ies>sub-if>grp-if>icmp`		
`icmp`	Configures ICMP parameters on an IES or VPLS IP interface.	749
`mask-reply`	Enables responses to ICMP mask requests on the router interface.	749
`redirects`	Enables and configures the rate for ICMP redirect messages issued on the router interface.	749
`ttl-expired`	Configures the rate ICMP TTL expired messages are issued by the IP interface.	750
`unreachables`	Enables and configures the rate for ICMP host and network destination unreachable messages issued on the router interface.	750
Configure VRRP parameters		698
`config>service>ies>interface`		
`authentication-key`	Assigns a simple text password authentication key to generate master VRRP advertisement messages and validate received VRRP advertisement messages.	795
`authentication-type`	Assigns the authentication method to generate master VRRP advertisement messages and validate received VRRP advertisement messages.	796
`backup`	Configures virtual router IP addresses for the interface.	798
`mac`	Assigns a specific MAC address to an IES or VPLS IP interface.	798
`master-int-inherit`	Allows the master instance to dictate the master down timer (non-owner context only).	798
`message-interval`	Sets the advertisement timer and indirectly sets the master down timer on the virtual router instance.	799
`ping-reply`	Enables the non-owner master to reply to ICMP echo requests directed at the virtual router instances IP addresses.	799
`policy`	Creates VRRP control policies.	800
`preempt`	Provides the ability of overriding an existing non-owner master to the virtual router instance.	800

Table 1: CLI Commands to Configure IES Service Parameters (Continued)

Command	Description	Page
priority	Provides the ability to configure a specific priority value to the virtual router instance.	801
no shutdown	Administratively disables VRRP.	717
ssh-reply	Enables the non-owner master to reply to SSH requests directed at the virtual router instances IP addresses.	801
telnet-reply	Enables the non-owner master to reply to TCP port 23 Telnet requests directed at the virtual router instances IP addresses.	802
traceroute-reply	Allows non-owner master to reply to traceroute requests (non-owner context only).	802
Configure IES interface IPv6 parameters		
config>service>ies>interface		
ipv6	Enables the context to configure IPv6 for an IES interface.	753
address	Assigns an IPv6 address to the IES interface.	753
dhcp6-relay	Enables the context to configure DHCPv6 relay parameters	753
lease-populate	Specifies the maximum number of DHCP6 lease states allocated by the DHCP6 DHCP relay function, allowed on this interface.	754
neighbor-resolution	Enables neighbor resolution with DHCPv6 relay.	754
option	Configures the DHCPv6 Relay information options	754
interface-id	Enables the sending of the interface-id option in the DHCPv6 relay packet.	754
remote-id	Enables the sending of the remote-id option in the DHCPv6 relay packet.	755
server	Configures the DHCPv6 server IPv6 address.	755
source-address	Configure the source IPv6 address of the DHCPv6 relay messages.	756
dhcp6-server	Configures DHCPv6 server parameters for the IES interface.	756
max-nbr-of-leases	Specifies the number of DHCPv6 leases allowed.	756
prefix-delegation	Administratively enables prefix delegation on this interface.	756
icmp6	Configures ICMPv6 parameters for the IES interface.	758
packet-too-big	Configures ICMPv6 packet-too-big messages.	758
param-problem	Configures ICMPv6 param-problem messages.	759
redirects	Configures ICMPv6 redirect messages.	759
time-exceeded	Configures ICMPv6 time-exceeded messages.	760
unreachables	Configures ICMPv6 unreachable messages.	760
neighbor	Configures IPv6-to-MAC address mapping on the IES interface.	761

Table 1: CLI Commands to Configure IES Service Parameters (Continued)

Command	Description	Page
Configure IES interface spoke-SDP parameters		
`config>service>ies>if# spoke-sdp`		
`spoke-sdp`	Binds a service to an existing SDP.	762
`egress`	Enables the context to specify the egress filter policy and VC label value.	762
`ingress`	Enables the context to specify the ingress filter policy and VC label value.	763
`filter`	Associates an IP filter policy with an egress SAP or IP interface.	775
`vc-label`	Configures the VC label.	763
Configure IES redundant interface parameters		
`config>service>ies`		
`redundant-interface`	Configures a subscriber interface.	721
`address`	Assigns an IP address/IP subnet format and a remote IP to the interface.	721
`description`	A text string that describes the redudant interface.	718
`no shutdown`	Administratively enables the interface.	717
Configure SAP parameters		695
`config>service>ies>interface`		
`config>service>ies>sub-if>grp-if`		
`sap` *`sap-id`*	Creates a SAP within the IES service.	764
`accounting-policy`	Configures the accounting policy that applies to the SAP.	774
`anti-spoof`	Enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP.	732
`atm`	Enables access to the context to configure ATM-related attributes only in the **config>service>ies>interface** context.	792
`collect-stats`	Enables the collection of accounting and statistical data for the SAP or network port.	774
`egress`	Configures egress SAP QoS policies and filter policies.	776
`host`	Creates a static subscriber host for the SAP.	740
`ingress`	Configures ingress SAP QoS policies and filter policies.	776
`multiservice-site`	Specifies the multi-service-site to which this SAP belongs.	767
`no shutdown`	Administratively enables the SAP.	717
Configure SAP ATM parameters		
`config>service>ies>if>sap>atm`		
`atm`	Enables access to the context to configure ATM-related attributes.	792
`egress`	Configures egress ATM attributes for the SAP.	792

Table 1: CLI Commands to Configure IES Service Parameters (Continued)

Command	Description	Page
`traffic-desc`	Assigns an ATM traffic descriptor profile to a given context (for example a SAP).	793
`encapsulation`	Specifies the data encapsulation for an ATM PVCC delimited SAP.	792
`ingress`	Configures ingress ATM attributes for the SAP.	793
`oam`	Enables the context to configure OAM functionality for a PVCC delimiting a SAP.	793
`alarm-cells`	Configures AIS/RDI fault management on a PVCC.	794
`periodic-loopback`	Enables periodic OAM loopbacks on this SAP.	794

Basic Configuration

The most basic IES service configuration has the following entities:

- Customer ID (refer to Configuring Customers on page 64)
- 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-A.

```
*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 2/1/5:0 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 a subscriber interface (optional).
5. Create an interface.
6. Define SAP parameters on the interface
   - → Select node(s) and port(s).
   - → Optional - select QoS policies other than the default (configured in the `config>qos` context).
   - → Optional - select filter policies (configured in the `config>filter` context).
   - → Optional - select accounting policy (configured in the `config>log` context).
7. Enable service.

Configuring IES Components

Use the CLI syntax to configure the following entities:

- Configuring an IES Service on page 691
- Configuring IES Interface Parameters on page 693
  - → Configuring IES Subscriber Interface Parameters on page 692
  - → Configuring Spoke-SDP Parameters on page 694
  - → Configuring SAP Parameters on page 695
  - → Configuring VRRP on page 698

Configuring an IES Service

Use the following CLI syntax to create an IES service:

CLI Syntax:
```
config>service# ies service-id customer customer-id [vpn vpn-id]
        description description-string
        subscriber-subnet ip-address/mask [igp-advertise {subnet|host-
         routes}]
        interface ip-int-name (see Configuring IES Interface Parameters on page 693)
        no shutdown
```

Example:
```
config>service# ies 1001 customer 1730 create
config>service>ies# no shutdown
config>service>ies# description "to-internet"
config>service>ies#
```

The following example displays the IES service creation output.

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

Configuring IES Subscriber Interface Parameters

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

Refer to IES Services Command Reference on page 705 for CLI syntax to configure IES subscriber interface parameters.

The following displays the command usage to configure IES interface parameters:

Example:

```
config>service# ies 1000
config>service>ies$ subscriber-interface "test1" create
config>service>ies>if$ address 143.144.145.1/24
config>service>ies>if# group-interface "new-if" create
config>service>ies>sub-if>grp-if# sap 1/2/19:0 create
config>service>ies>sub-if>grp-if>sap# ingress
config>service>ies>sub-if>grp-if>sap>ingress# filter ip 10
config>service>ies>sub-if>grp-if>sap>ingress# qos 2
config>service>ies>sub-if>grp-if>sap>ingress# exit
```

The following example displays the subscriber interface configuration:

```
A:ALA-48>config>service>ies>sub-if# info
----------------------------------------------
                    address 143.144.145.1/24
                    group-interface "new-if" create
                        sap 1/2/19:0 create
                            ingress
                                qos 2
                                filter ip 10
                            exit
                          host ip 143.144.145.100 mac 00:01:00:00:00:01
                        exit
                    exit
----------------------------------------------
A:ALA-48>config>service>ies>sub-if#
```

Configuring IES Interface Parameters

Refer to IES Services Command Reference on page 705 for CLI syntax.

The following displays the command usage to configure IES interface parameters:

Example:

```
config>service# ies 1000
config>service>ies$ interface to-web create
config>service>ies>if$ address 10.1.1.1/24
config>service>ies>if# no shutdown
```

The following example displays the IES configuration:

```
A:ALA-48>config>service>ies>if# info
----------------------------------------------
               address 10.1.1.1/24
               sap 2/1/10:50 create
                   ingress
                       qos 100
                   exit
                   egress
                       scheduler-policy "SLA1"
                   exit
               exit
               vrrp 1 owner
                   authentication-type password
                   authentication-key "3WErEDozxyQ" hash
               exit
----------------------------------------------
A:ALA-48>config>service>ies>if#
```

Configuring Spoke-SDP Parameters

Use the following CLI syntax to configure spoke SDP parameters:

CLI Syntax: config>service# ies *service-id* [customer *customer-id*] [vpn *vpn-id*]

```
interface ip-int-name
   spoke-sdp sdp-id:vc-id
      egress
         filter {ip ip-filter-id}
         vc-label egress-vc-label
      ingress
         filter {ip ip-filter-id}
         vc-label ingress-vc-label
      no shutdown
```

Example:

```
config>service# ies 1001 customer 1730 create
config>service>ies$ description "to internet"
config>service>ies# interface "spokeSDP-test" create
config>service>ies>if$ spoke-sdp 2:100 create
config>service>ies>if>spoke-sdp$ egress
config>service>ies>if>spoke-sdp>egress$ filter ip 10
config>service>ies>if>spoke-sdp>egress$ exit
config>service>ies>if>spoke-sdp# no shutdown
```

The following example displays the spoke SDP configuration.

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

Configuring SAP Parameters

A SAP is a combination of a port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique within a router.

When configuring IES SAP parameters, a default QoS policy is applied to each ingress and egress SAP. Additional QoS policies and scheduler policies must be configured in the `config>qos` context. Filter policies are configured in the `config>filter` context and must be explicitly applied to a SAP. There are no default filter policies.

IES interface ATM SAP parameters can only be configured on ATM-type MDAs and ATM-configured ports. See the *7750 SR OS Basic System Configuration Guide.*

Refer to IES Services Command Reference on page 705 for CLI syntax.

The following displays the command usage to configure IES interface parameters:

Example:

```
config>service# ies 1000
config>service>ies# interface test
config>service>ies>if# sap 5/1/3.1:0 create
config>service>ies>if>sap# ingress
config>service>ies>if>sap>ingress# qos 101
config>service>ies>if>sap>ingress# exit
config>service>ies>if>sap# egress
config>service>ies>if>sap>egress# scheduler-policy alpha
config>service>ies>if>sap>egress# qos 1010
config>service>ies>if>sap>egress# exit
config>service>ies>if>sap# exit
config>service>ies>if#
```

This example displays an IES SAP configuration.

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

Configuring IES SAP ATM Parameters

Use the following CLI syntax to configure IES SAP ATM parameters.

CLI Syntax:
```
config>service# ies service-id
        interface interface-name
           sap sap-id
              atm
                 egress
                    traffic-desc traffic-desc-profile-id
                 encapsulation encapsulation-type
                 ingress
                    traffic-desc traffic-desc-profile-id
                 oam
                    alarm-cells
                    terminate
              collect-stats
              description description-string
              egress
                 qos policy-id
                 scheduler-policy scheduler-policy-name
              ingress
                 qos policy-id [shared-queuing]
                 scheduler-policy scheduler-policy-name
              multi-service-site customer-site-name
              no shutdown
```

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

PE router 1 (A:ALA-41):

Example:
```
A:ALA-41>config>service# ies 5
A:ALA-41>config>service>ies# sap 1/1/1:0/32 create
A:ALA-41>config>service>ies>sap# ingress
A:ALA-41>config>service>ies>sap>ingress# qos 102
A:ALA-41>config>service>ies>sap>ingress# exit
A:ALA-41>config>service>ies>sap# egress
A:ALA-41>config>service>ies>sap>egress# qos 103
A:ALA-41>config>service>ies>sap>egress# exit
A:ALA-41>config>service>ies>sap# no shutdown
A:ALA-41>config>service>ies>sap# exit
A:ALA-41>config>service>ies#
```

PE router 2 (A:ALA-42):

Example:
```
A:ALA-42>config>service# ies 5
A:ALA-42>config>service>ies# sap 2/2/2:0/32 create
A:ALA-42>config>service>ies>sap# ingress
A:ALA-42>config>service>ies>sap>ingress# qos 102
A:ALA-42>config>service>ies>sap>ingress# exit
```

```
A:ALA-42>config>service>ies>sap# egress
A:ALA-42>config>service>ies>sap>egress# qos 103
A:ALA-42>config>service>ies>sap>egress# exit
A:ALA-42>config>service>ies>sap# no shutdown
A:ALA-42>config>service>ies>sap# exit
A:ALA-42>config>service>ies#
```

The following output displays the IES SAP configuration.

PE Router 1 (ALA-41):

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

Configuring VRRP

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

For further information about VRRP CLI syntax and command descriptions refer to the *7750 SR OS Router Configuration Guide*.

Use the following CLI syntax to configure non-owner VRRP IES interface parameters.

CLI Syntax:
```
config>service# ies service-id
         interface ip-int-name
            vrrp vrid
               authentication-key [authentication-key | hash-key] [hash]
               authentication-type {password | message-digest}
               backup ip-addr
               mac ieee-mac-address
               master-int-inherit
               priority
               policy (vrrp instance) vrrp-policy-id
               preempt
               message-interval seconds
               ping-reply
               ssh-reply
               telnet-reply
               traceroute-reply
               no shutdown
```

Use the following CLI syntax to configure owner VRRP IES interface parameters.

CLI Syntax:
```
config>service> ies service-id
         interface ip-int-name
            vrrp vrid owner
               authentication-type {password | message-digest}
               authentication-key [authentication-key | hash-key] [hash]
               backup ip-addr
               mac ieee-mac-address
               message-interval seconds
```

The following displays the command usage to configure owner VRRP interface parameters:

Example:
```
config>service# ies 1000
config>service>ies# interface test
config>service>ies>if$ address 10.10.36.2/24
config>service>ies>if$ vrrp 2 owner
config>service>ies>if>vrrp$ backup 10.10.36.2
config>service>ies>if>vrrp$ authentication-type password
config>service>ies>if>vrrp$ authentication-key alcatel
config>service>ies>if>vrrp$
```

The following example displays the IES configuration:

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

Service Management Tasks

This section discusses the following service management tasks:

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 charges 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 command usage to modify IES service parameters:

Example:

```
config>service# ies 1000
config>service>ies# description "This is a new description"
config>service>ies# interface "to-web"
config>service>ies>if# allow-directed-broadcast
config>service>ies>if# icmp
config>service>ies>if# mac  00:dc:98:1d:00:00
config>service>ies>if# no shutdown
config>service>ies>if# exit
config>service>ies# no shutdown
config>service>ies# exit
```

```
*A:ALA-A>config>service>ies# info
----------------------------------------------
    ies 1000 customer 50 vpn 1000 create
          description "This is a new description"
          interface "to-web" create
              address 10.1.1.1/24
              mac 00:dc:98:1d:00:00
              allow-directed-broadcast
              sap 2/1/50:0 create
              exit
          exit
          no shutdown
    exit
----------------------------------------------
*A:ALA-A>config>service#
```

Deleting a Spoke-SDP

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

Use the following CLI syntax to delete a spoke SDP from an interface:

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

Example:
```
config>service# ies 1001
config>service>ies# interface "spokeSDP-test"
config>service>ies>if# spoke-sdp 2:100
config>service>ies>if>spoke-sdp# shutdown
config>service>ies>if>spoke-sdp# exit
config>service>ies>if# no spoke-sdp 2:100
config>service>ies>if# exit
```

The following example displays the spoke SDP configuration.

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

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
```

Example:
```
config>service# ies 1000
config>service>ies# shutdown
config>service>ies# interface "to-web"
config>service>ies>if# shutdown
config>service>ies>if# sap 2/1/50:0
config>service>ies>if>sap# shutdown
config>service>ies>if>sap# exit
config>service>ies>if# no sap 2/1/50:0
config>service>ies>if# exit
config>service>ies# no interface "to-web"
config>service>ies# exit
config>service# no ies 1000
```

The service is removed from the configuration.

Disabling an IES Service

You can shut down an IES service without deleting the service parameters.

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

Example:
```
config>service# ies 2000
config>service>ies# shutdown
config>service>ies# exit
```

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

config
- **service**
  - **ies** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*]
    - **description** *description-string*
    - **no description**
    - [**no**] **shutdown**
    - [**no**] **interface** *ip-int-name*
      - [**no**] **active-cpm-protocols**
      - **address** *ip-address/mask* [**netmask**] [**broadcast** {**all-ones** | **host-ones**}]
      - **no address**
      - [**no**] **allow-directed-broadcasts**
      - [**no**] **arp-populate**
      - **arp-timeout** *seconds*
      - **no arp-timeout**
      - **authentication-policy** *name name*
      - **no authentication-policy** *name*
      - **bfd** *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier*]
      - **no bfd**
      - **cflowd** [**acl** | **interface**]
      - **no cflowd**
      - **description** *description-string*
      - **no description**
      - **dhcp**
        - **description** *description-string*
        - **no description**
        - **gi-address** *ip-address* [*src-ip-addr*]
        - **no gi-address**
        - **lease-populate** [*nbr-of-leases*]
        - **no lease-populate**
        - [**no**] **option**
          - **action** {**replace** | **drop** | **keep**}

- **no action**
- **circuit-id** [**ascii-tuple** | **ifindex** | **sap-id** |**vlan-ascii-tuple**]
- **no circuit-id**
- **remote-id** [**mac** | **string** *string*]
- **no remote-id**
- [**no**] **vendor-specific-option**
  - [**no**] **client-mac-address**
  - [**no**] **sap-id**
  - [**no**] **service-id**
  - **string** ***text***
  - **no string**
  - [**no**] **system-id**

- **proxy-server**
  - **emulated-server** ***ip-address***
  - **no emulated-server**
  - **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*] [**radius-override**]
  - **no lease-time**
  - [**no**] **shutdown**
- **server** *server1* [*server2*...(up to 8 max)]
- **no server**
- [**no**] **shutdown**
- [**no**] **trusted**

- **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**
- **ip-mtu** *octets*
- **no ip-mtu**
- [**no**] **ipv6**
  - **address** *ipv6-address/prefix-length* [**eui-64**]
  - **no address** *ipv6-address/prefix-length*
  - [**no**] **dhcp6-relay**
    - **description** *description-string*
    - **no description**
    - **lease-populate** [*nbr-of-leases*]
    - **no lease-populate**
    - [**no**] **neighbor-resolution**
    - [**no**] **option**
      - **interface-id**
      - **interface-id ascii-tuple**
      - **interface-id ifindex**
      - **interface-id sap-id**
      - **interface-id string**
      - **no interface-id**
      - [**no**] **remote-id**
    - **server** *ipv6z-address* [*ipv6z-address*...(up to 8 max)]
    - **no server** [*ipv6z-address*...(up to 8 max)]
    - [**no**] **shutdown**
    - **source-address** *ipv6-address*

```
                    — no source-address
                — [no] dhcp6-server
                    — max-nbr-of-leases max-nbr-of-leases
                    — no max-nbr-of-leases
                    — [no] prefix-delegation
                        — [no] prefix ipv6-address/prefix-length
                            — duid duid [iaid iaid]
                            — no duid
                            — preferred-lifetime seconds
                            — preferred-lifetime infinite
                            — no preferred-lifetime
                            — valid-lifetime seconds
                            — valid-lifetime infinite
                            — no valid-lifetime
                        — [no] shutdown
                — icmp6
                    — packet-too-big [number seconds]
                    — no packet-too-big
                    — param-problem [number seconds]
                    — no param-problem
                    — redirects [number seconds]
                    — no redirects
                    — time-exceeded [number seconds]
                    — no time-exceeded
                    — unreachables [number seconds]
                    — no unreachables
                — [no] local-proxy-nd
                — neighbor ipv6-address mac-address
                — no neighbor ipv6-address
                — proxy-nd-policy policy-name [policy-name...(up to 5 max)]
                — no proxy-nd-policy
            — [no] local-proxy-arp
            — [no] loopback
            — [no] mac ieee-address
            — [no] proxy-arp-policy policy-name [policy-name...(up to 5 max)]
            — [no] remote-proxy-arp
    — [no] redundant-interface ip-int-name
            — address {ip-address/mask | ip-address netmask} [remote-ip ip-address]
            — no address
            — description description-string
            — no description
            — [no] shutdown
            — no spoke-sdp sdp-id:vc-id
                — egress
                    — filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                    — vc-label egress-vc-label
                    — no vc-label [egress-vc-label]
                — ingress
                    — filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                    — vc-label ingress-vc-label
                    — no vc-label [ingress-vc-label]
            — [no] shutdown
```

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id]
            — [no] interface ip-int-name
                — [no] sap sap-id
                    — accounting-policy acct-policy-id
                    — no accounting-policy [acct-policy-id]
                    — anti-spoof {ip | ip-mac}
                    — no anti-spoof
                    — atm
                        — egress
                            — traffic-desc traffic-desc-profile-id
                            — no traffic-desc
                        — encapsulation atm-encap-type
                        — ingress
                            — traffic-desc traffic-desc-profile-id
                            — no traffic-desc
                        — oam
                            — [no] alarm-cells
                            — [no] periodic-loopback
                    — [no] collect-stats
                    — description description-string
                    — no description
                    — egress
                        — agg-rate-limit agg-rate
                        — no agg-rate-limit
                        — [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]
                        — qos sap-egress-policy-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
                                — rate pir-rate [cir cir-rate]
                                — no rate
                        — [no] scheduler-override
                            — [no] scheduler scheduler-name
                                — rate pir-rate [cir cir-rate]
                                — no rate
                        — scheduler-policy scheduler-policy-name
                        — no scheduler-policy
                    — host {[ip ip-address] [mac ieee-address]}[subscriber sub-
                      ident-string] [sub-profile sub-profile-name] [sla-
                      profile sla-profile-name]
```

```
— no host {[ip ip-address] [mac ieee-address]}
— no host all
— ingress
    — filter [ip ip-filter-id]
    — filter [ipv6 ipv6-filter-id]
    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
    — match-qinq-dot1p {top|bottom}
    — no match-qinq-dot1p
    — qos policy-id [shared-queuing | multipoint-shared]
    — 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
            — rate pir-rate [cir cir-rate]
            — no rate
    — [no] scheduler-override
        — [no] scheduler scheduler-name
            — rate pir-rate [cir cir-rate]
            — no rate
    — scheduler-policy scheduler-policy-name
    — no scheduler-policy
— multi-service-site customer-site-name
— no multi-service-site
— tod-suite tod-suite-name
— no tod-suite
— [no] shutdown
— secondary {ip-address/mask | ip-address netmask } [broadcast all-ones |
  host-ones] [igp-inhibit]
— no secondary ip-address
— [no] shutdown
— [no] static-arp ip-address
— tos-marking-state {trusted | untrusted}
— no tos-marking-state
— unnumbered [ip-int-name | ip-address]
— no unnumbered
```

VRRP Commands

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id]
            — [no] interface ip-int-name
                — vrrp virtual-router-id [owner]
                — no vrrp virtual-router-id
```

- **authentication-key** {*authentication-key* | *hash-key*} [**hash** | **hash2**]
- **no authentication-key**
- **authentication-type** {**password** | **message-digest}**
- **no authentication-type**
- [**no**] **backup** *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*
      - **spoke-sdp** *sdp-id:vc-id*
      - **no spoke-sdp** *sdp-id:vc-id*
        - **egress**
          - **filter** [**ip** *ip-filter-id*]
          - **filter** [**ipv6** *ipv6-filter-id*]
          - **no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
          - **vc-label** *egress-vc-label*
          - **no vc-label** [*egress-vc-label*]
        - **ingress**
          - **filter** {**ip** *ip-filter-id*}
          - **no filter**
          - **vc-label** *ingress-vc-label*
          - **no vc-label** [*ingress-vc-label*]

Subscriber Interface Commands

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id]
            — [no] subscriber-interface ip-int-name
                — [no] address {ip-address/mask | ip-address netmask} [gw-ip-address ip-address]
                — description description-string
                — no description
                — dhcp
                    — gi-address ip-address [src-ip-addr]
                    — no gi-address
                — [no] group-interface ip-int-name
                    — [no] arp-populate
                    — arp-timeout seconds
                    — no arp-timeout
                    — authentication-policy name
                    — no authentication-policy
                    — description description-string
                    — no description
                    — dhcp
                        — description description-string
                        — no description
                        — gi-address ip-address [src-ip-addr]
                        — no gi-address
                        — lease-populate nbr-of-leases
                        — no lease-populate
                        — [no] match-circuit-id
                        — [no] option
                            — action {replace | drop | keep}
                            — no action
                            — circuit-id [ascii-tuple | ifindex | sap-id |vlan-ascii-tuple]
                            — no circuit-id
                            — remote-id [mac | string string]
                            — no remote-id
                            — [no] vendor-specific-option
                                — [no] client-mac-address
                                — [no] sap-id
                                — [no] service-id
                                — string text
                                — no string
                                — [no] system-id
                        — proxy-server
                            — emulated-server ip-address
                            — no emulated-server
                            — lease-time [days days] [hrs hours] [min minutes] [sec seconds] [radius-override]
                            — no lease-time
                            — [no] lease-time
                            — [no] shutdown
                        — server server1 [server2...(up to 8 max)]
```

- **no server**
- [**no**] **shutdown**
- [**no**] **trusted**

— **host-connectivity-verify** [**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**

— [**no**] **local-proxy-arp**
— [**no**] **mac** *ieee-address*
— [**no**] **proxy-arp-policy** *policy-name* [*policy-name*...(up to 5 max)]
— **redundant-interface** *red-ip-int-name*
— **no redundant-interface**
— [**no**] **remote-proxy-arp**
— [**no**] **sap** *sap-id*
- **accounting-policy** *acct-policy-id*
- **no accounting-policy** [*acct-policy-id*]
- **anti-spoof** {**ip** | **ip-mac**}
- **no anti-spoof**
- **atm**
  - **egress**
    - **traffic-desc** *traffic-desc-profile-id*
    - **no traffic-desc**
  - **encapsulation** ***atm-encap-type***
  - **ingress**
    - **traffic-desc** *traffic-desc-profile-id*
    - **no traffic-desc**
  - **oam**
    - [**no**] **alarm-cells**
    - [**no**] **periodic-loopback**
- [**no**] **collect-stats**
- **description** *description-string*
- **no description**
- **egress**
  - **agg-rate-limit** ***agg-rate***
  - **no agg-rate-limit** ***agg-rate***
  - **filter ip** *ip-filter-id*
  - **no filter**
  - **filter ipv6** ***ipv6-filter-id***
  - **no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
  - [**no**] **qinq-mark-top-only**
  - **qos** ***policy-id***
  - **no qos** ***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 | default}
                        — no mbs
                        — rate  pir-rate [cir cir-rate]
                        — no rate
                    — scheduler-policy scheduler-policy-name
                    — scheduler-policy
                — host ip ip-address [mac ieee-address] [subscriber
                  sub-ident-string] [sub-profile sub-profile-name] [sla-
                  profile sla-profile-name] [ancp-string ancp-string]
                — no host {[ip ip-address] [mac ieee-address]}
                — no host all
                — ingress
                    — filter ip ip-filter-id
                    — no filter
                    — filter ipv6 ipv6-filter-id
                    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                    — match-qinq-dot1p {top|bottom}
                    — no match-qinq-dot1p
                    — qos policy-id [shared-queuing]
                    — no qos policy-id
                    — scheduler-policy scheduler-policy-name
                    — no scheduler-policy
                — multi-service-site customer-site-name
                — no multi-service-site
                — [no] shutdown
                — [no] sub-sla-mgmt
                    — def-sla-profile  default-sla-profile-name
                    — no def-sla-profile
                    — def-sub-profile  default-subscriber-profile-name
                    — no def-sub-profile
                    — multi-sub-sap subscriber-limit
                    — no multi-sub-sap
                    — [no] shutdown
                    — single-sub-parameters
                        — non-sub-traffic sub-profile sub-profile-
                          name  sla-profile sla-profile-name [sub-
                          scriber sub-ident-string]
                        — no non-sub-traffic
                        — [no] profiled-traffic-only
                    — sub-ident-policy sub-ident-policy-name
                    — no sub-ident-policy
                    — tod-suite tod-suite-name
                    — no tod-suite
                — [no] shutdown
        — [no] srrp srrp-id
                — description description-string
                — no description
                — gw-mac mac-address
                — no gw-mac
                — keep-alive-interval interval
                — no keep-alive-interval
                — message-path sap-id
```

— **no message-path**
— [**no**] **policy** *vrrp-policy-id*
— **priority** *priority*
— **no priority**
— [**no**] **shutdown**

Show Commands

show
- **service**
  - **customer** [*customer-id*] [**site** *customer-site-name*]
  - **egress-label** *start-label* [*end-label*]
  - **ingress-label** *end-label* [*end-label*]
  - **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*
  - **service-using** [**ies**] [**customer** *customer-id*]
  - **subscriber-using** [**service-id** *service-id*] [**sap-id** *sap-id*] [**interface** *ip-int-name*] [**ip** *ip-address[/mask]*] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*]
  - **id** *service-id*
    - **all**
    - **arp** [*ip-address*]|[**mac** *ieee-address*]|[**sap** *sap-id*]|[**interface** *ip-int-name*]|[**sdp** *sdp-id:vc-id*]
    - **authentication**
      - **statistics** [**policy** *name*] [**sap** *sap-id*]
    - **base**
    - **dhcp**
      - **lease-state** [[**sap** *sap-id*] | [**sdp** *sdp-id*:*vc-id*] | [**interface** *interface-name*] | [**ip-address** *ip-address*]] [**detail**]
      - **statistics** [**sap** *sap-id*
      - **statistics** [**sdp** *sdp-id*:*vc-id*]
      - **statistics** [**interface** *interface-name*]
      - **summary**
    - **gsmp**
      - **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*]
    - **labels**
    - **retailers**
    - **sap** *sap-id* [**detail**]]
    - **sdp** [{*sdp-id* | **far-end** *ip-address*}] [**detail**]]
    - **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-address*[/*mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**detail**]
    - **wholesalers**
- **router**
  - **dhcp**
    - **statistics** [*ip-int-name* | *ip-address*]
    - **summary**

Clear Commands

clear
- **router**
  - **dhcp**
    - **statistics** [*ip-int-name* | *ip-address*]
  - **interface** [*ip-int-name* | *ip-address*] [**icmp**]

clear
- **service**
  - **id** *service-id*
    - **fdb** {**all** | **mac** *ieee-address* | **sap** *sap-id* | **mesh-sdp** *sdp-id*[**:***vc-id*] | **spoke-sdp** *sdp-id***:***vc-id*}
    - **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*
    - **dhcp6**
      - **lease-state** [**ip-address** *ipv6-address*/*prefix-length*] [**mac** *ieee-address*]
      - **statistics** [**interface** *ip-int-name* | *ipv6-address*]
    - **spoke-sdp** *sdp-id:vc-id ingress-vc-label*
    - **stp**

IES Service Configuration Commands

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>service>ies
config>service>ies>sub-if
config>service>ies>sub-if>grp-if
config>service>ies>sub-if>grp-if>dhcp
config>service>ies>sub-if>grp-if>sap
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>srrp
config>service>ies>interface
config>service>ies>if>vrrp
config>service>ies>if>dhcp
config>service>ies>if>dhcp>proxy-server
config>service>ies>redundant-interface

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

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

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

IES IP Interfaces — When the IP interface is shutdown, it will enter the administratively and operationally down states. For a SAP bound to the IP interface, no packets will be 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>sub-if
config>service>ies>sub-if>grp-if
config>service>ies>sub-if>grp-if>dhcp
config>service>ies>if>dhcp
config>service>ies>redundant-interface
config>service>ies>sub-if>grp-if>srrp

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

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

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

Default No description associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters 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*
ies *service-id*
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 Ethernet 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 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 7750 SR on which this service is defined.

Values 1 — 2147483647

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)

Redundant Interface Commands

redundant-interface

Syntax [**no**] **redundant-interface** *ip-int-name*

Context config>service>ies

Description This command configures a redundant interface.

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

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*} [**remote-ip** *ip-address*]
no address

Context config>service>ies>redundant-interface

Description This command assigns an IP address mask or netmask and a remote IP address to the interface.

Parameters *ip-address/mask* — Assigns an IP address/IP subnet format to the interface.

ip-address 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.

Assigns an IP address netmask to the interface.

remote-ip *ip-address* — Assigns a remote IP to the interface.

IES Subscriber Interface Commands

subscriber-interface

Syntax [**no**] **subscriber-interface** *ip-int-name*

Context config>service>ies

Description This command allows the operator to create special subscriber-based interfaces. It is used to contain multiple group interfaces. Multiple subnets associated with the subscriber interface can be applied to any of the contained group interfaces in any combination. The subscriber interface allows subnet sharing between group interfaces.

Use the **no** form of the command to remove the subscriber interface.

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

group-interface

Syntax [**no**] **group-interface** *ip-int-name*

Context config>service>ies>sub-if

Description This command enables the context to configure a group interface. A group interface is an interface that may contain one or more SAPs. This interface is used in triple-play services where multiple SAPs are part of the same subnet.

Default **none**

Parameters *ip-int-name* — Configures the interface group name. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

authentication-policy

Syntax **authentication-policy** *name*
no authentication-policy

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

Description This command assigns an authentication policy to the interface.

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

Default **no authentication-policy**

Parameters *name* — Specifies the authentication policy name. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

srrp

Syntax [**no**] **srrp** *srrp-id*

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

Description This command creates a Subscriber Router Redundancy Protocol (SRRP) instance on a group IP interface. An SRRP instance manages all subscriber subnets within the group interfaces subscriber IP interface or other subscriber IP interfaces that are associated through a wholesale/retail relationship. Only one unique SRRP instance can be configured per group interface.

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

Default **no srrp**

Parameters *srrp-id* — Specifies a 32 bit instance ID that must be unique to the system. The instance ID must also match the instance ID used by the remote router that is participating in the same SRRP context. SRRP is intended to perform a function similar to VRRP where adjacent IP hosts within local subnets use a default gateway to access IP hosts on other subnets.

Values 1 — 4294967295

gw-mac

Syntax **gw-mac** *mac-address*
no gw-mac

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

Description This command overrides the default SRRP gateway MAC address used by the SRRP instance. Unless specified, the system uses the same base MAC address for all SRRP instances with the last octet overridden by the lower 8 bits of the SRRP instance ID. . The same SRRP gateway MAC address should be in-use by both the local and remote routers participating in the same SRRP context.

One reason to change the default SRRP gateway MAC address is if two SRRP instances sharing the same broadcast domain are using the same SRRP gateway MAC. The system will use the SRRP instance ID to separate the SRRP messages (by ignoring the messages that does not match the local instance ID), but a unique SRRP gateway MAC is essential to separate the routed packets for each gateway IP address.

The **no** form of the command removes the explicit SRRP gateway MAC address from the SRRP instance. The SRRP gateway MAC address can only be changed or removed when the SRRP instance is shutdown.

Parameters *mac-address* — Specifies a MAC address that is used to override the default SRRP base MAC address

Values Any MAC address except all zeros, broadcast or multicast addresses. The offset is expressed in normal Ethernet MAC address notation. The defined gw-mac cannot be 00:00:00:00:00:00, ff:ff:ff:ff:ff:ff or any multicast address.

If not specified, the system uses the default SRRP gateway MAC address with the last octet set to the 8 least significant bits of the SRRP instance ID.

keep-alive-interval

Syntax **keep-alive-interval** *interval*
no keep-alive-interval

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

Description This command defines the interval between SRRP advertisement messages sent when operating in the master state. The interval is also the basis for setting the master-down timer used to determine when the master is no longer sending. The system uses three times the keep-alive interval to set the timer. Every time an SRRP advertisement is seen that is better then the local priority, the timer is reset. If the timer expires, the SRRP instance assumes that a master does not exist and initiates the attempt to become master.

When in backup state, the SRRP instance takes the keep-alive interval of the master as represented in the masters SRRP advertisement message. Once in master state, the SRRP instance uses its own configured keep-alive interval.

The keep-alive-interval may be changed at anytime, but will have no effect until the SRRP instance is in the master state.

The **no** form of the command restores the default interval.

Parameters *interval* — Specifies the interval between SRRP advertisement messages sent when operating in the master state.

Values 1 — 100 hundreds of milli-seconds

Default 1

message-path

Syntax **message-path** *sap-id*
no message-path

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

Description This command defines a specific SAP for SRRP in-band messaging. A message-path SAP must be defined prior to activating the SRRP instance. The defined SAP must exist on the SRRP instances group IP interface for the command to succeed and cannot currently be associated with any dynamic or static subscriber hosts. Once a group IP interface SAP has been defined as the transmission path for SRRP Advertisement messages, it cannot be administratively shutdown, will not support static or dynamic subscriber hosts and cannot be removed from the group IP interface.

The SRRP instance message-path command may be executed at anytime on the SRRP instance. Changing the message SAP will fail if a dynamic or static subscriber host is associated with the new SAP. Once successfully changed, the SRRP instance will immediately disable anti-spoof on the SAP and start sending SRRP Advertisement messages if the SRRP instance is activated.

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

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

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

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

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

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

policy

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

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

Description This command associates one or more VRRP policies with the SRRP instance. A VRRP policy is a collection of connectivity and verification tests used to manipulate the in-use priorities of VRRP and SRRP instances. A VRRP policy can test the link state of ports, ping IP hosts, discover the existence of routes in the routing table or the ability to reach L2 hosts. When one or more of these tests fail, the VRRP policy has the option of decrementing or setting an explicit value for the in-use priority of an SRRP instance.

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

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

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

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

Values 1 — 9999

priority

Syntax **priority** *priority*
no priority

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

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

The base priority of an SRRP instance can be managed by VRRP policies. A VRRP policy defines a set of connectivity or verification tests which, when they fail, may lower an SRRP instances base priority (creating an in-use priority for the instance). Every time an SRRP instances in-use priority changes when in master state, it sends an SRRP advertisement message with the new priority. If the dynamic priority drops to zero or receives an SRRP Advertisement message with a better priority, the SRRP instance transitions to the *becoming backup* state.

When the priority command is not specified, or the no priority command is executed, the system uses a default base priority of 100. The priority command may be executed at anytime.

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

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

Values 1 — 254

Default 100

IES Interface Commands

interface

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

Context config>service>ies

Description This command creates a logical IP routing interface for an Internet Ethernet 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 IP the interface and all the associated configuration. The interface must be administratively shutdown before issuing the **no interface** command.

For IES services, the IP interface must be shutdown before the SAP on that interface may be removed. 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.

active-cpm-protocols

Syntax [**no**] **active-cpm-protocols**

Context config>service>ies>interface

Description This command enables CPM protocols on this interface.

address

Syntax **address** *ip-address/mask-length*
address *ip-address mask*
address *ip-address*
no address

Context config>service>ies>interface
config>service>ies>subscriber-interface

Description This command assigns an IP address, IP subnet, and broadcast address format to an IES IP router interface. Only one IP address can be associated with an IP interface.

An IP address must be assigned to each IES IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the 7750 SR.

In the IES subscriber interface context, this command is used to assign one or more (16 maximum) host IP addresses and subnets. This differs from a normal IES interfaces where the **secondary** command creates an additional subnet after the primary address is assigned. A user can then add or remove addresses without having to keep a primary address.

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

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

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

Use the **no** form of this command to remove the IP address assignment from the IP interface. When the **no address** command is entered, the interface becomes operationally down.

Address	Admin state	Oper state
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an adminstratively 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 /30 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.

address

Syntax [**no**] **address** {*ip-address/mask* | *ip-address netmask*} [**gw-ip-address** *ip-address*]

Context config>service>ies>sub-if

Description This command configures the local subscriber subnets available on a subscriber IP interface. The configured ip-address and mask define the address space associated with the subscriber subnet. Up to 16 IP subnets can be created on a single subscriber IP interface. Each subnet supports a locally owned IP host address within the subnet that is not expected to appear on other routers that may be servicing the same subscriber subnet. For redundancy purposes, the keyword **gw-address** defines a separate IP address within the subnet for Subscriber Routed Redundancy Protocol (SRRP) routing. This IP address must be the same on the local and remote routers participating in a common SRRP instance.

In SRRP, a single SRRP instance is tied to a group IP interface. The group IP interface is contained directly within a subscriber IP interface context and thus directly associated with the subscriber subnets on the subscriber IP interface. The SRRP instance is also indirectly associated with any subscriber subnets tied to the subscriber interface through wholesale/retail VPRN configurations. With the directly-associated and the indirectly-associated subscriber interface subnets, a single SRRP instance can manage hundreds of SRRP gateway IP addresses. This automatic subnet association to the SRRP instance is different from VRRP where the redundant IP address is defined within the VRRP context.

Defining an SRRP gateway IP address on a subscriber subnet is not optional when the subnet is associated with a group IP interface with SRRP enabled. Enabling SRRP (**no shutdown**) will fail if one or more subscriber subnets do not have an SRRP gateway IP address defined. Creating a new subscriber subnet without an SRRP gateway IP address defined will fail when the subscriber subnet is associated with a group IP interface with an active SRRP instance. Once SRRP is enabled on a group interface, the SRRP instance will manage the ARP response and routing behavior for all subscriber hosts reachable through the group IP interface.

The no form of the command removes the address from a subscriber subnet. The **address** command for the specific subscriber subnet must be executed without the **gw-address** parameter. To succeed, all SRRP instances associated with the subscriber subnet must removed or shutdown.

Parameters *ip-address/mask* | *ip-address netmask* — Specifies the address space associated with the subscriber subnet

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

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

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

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

allow-directed-broadcasts

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

Context config>service>ies>interface

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

A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address 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** or **ip-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.

anti-spoof

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

Context config>service>ies>sub-if>grp-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**) 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 **ip-mac**

Parameters **ip —** Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the anti-spoof type **ip** command will fail.

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.

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

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

Description This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

When the **arp-populate** and **lease-populate** commands are enabled on an IES interface, the ARP table entries will no longer be dynamically learned, but instead by snooping DHCP ACK message from a DHCP server. In this case the configured **arp-timeout** value has no effect.

The default value for **arp-timeout** is 14400 seconds (4 hours).

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

bfd

Syntax **bfd** *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier*]
no bfd

Context config>service>ies>interface

Description This command specifies the BFD parameters for the associated IP interface. If no parameters are defined the default value are used.

The multiplier specifies the number of consecutive BFD messages that must be missed from the peer before the BFD session state is changed to down and the upper level protocols (OSPF, IS-IS or PIM) is notified of the fault.

The **no** form of the command removes BFD from the associated IGP protocol adjacency.

Default **no bfd**

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

Values 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 2 — 20

Default 3

cflowd

Syntax **cflowd** {**acl** | **interface**}
no cflowd

Context config>service>ies>interface

Description This command enables **cflowd** to collect traffic flow samples through a router for analysis.

cflowd is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement. When **cflowd** is enabled at the interface level, all packets forwarded by the interface are subjected to analysis according to the **cflowd** configuration.

Default **no cflowd**

Parameters *ACL — cflowd* configuration associated with a filter.

interface — cflowd configuration associated with an IP interface.

local-proxy-arp

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

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

Description This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus will become the forwarding point for all traffic between hosts in that subnet.

When local-proxy-arp is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

Default **ies>interface: no local-proxy-arp**

ies>sub-if>grp-if: local-proxy-arp

ip-mtu

Syntax **ip-mtu** *octets*
no ip-mtu

Context config>service>ies>interface

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**

loopback

Syntax [**no**] **loopback**

Context config>service>ies>interface

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** *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.

Default **none**

mac

Syntax **mac** *ieee-address*
no mac

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

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

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

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

Default The physical MAC address associated with the Ethernet interface that the SAP is configured on (the default MAC address assigned to the interface, assigned by the system).

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

secondary

Syntax **secondary** {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**] [**igp-inhibit**]
no secondary *ip-address*

Context config>service>ies>interface

Description This command assigns a secondary IP address/IP subnet/broadcast address format to the interface.

Default none

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 /30 subnets).

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical 'AND' function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. 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 indictates 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.

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** *ip-address ieee-mac-address*
no static-arp *ip-address*

Context config>service>ies>interface

Description This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP 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.

tos-marking-state

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

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

Description This command configures the interface as an unnumbered interface.

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.

proxy-arp-policy

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

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

remote-proxy-arp

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

Description This command enables remote proxy ARP on the interface.

Remote proxy ARP is similar to proxy ARP. It allows the router to answer an ARP request on an interface for a subnet that is not provisioned on that interface. This allows the router to forward to the other subnet on behalf of the requester. To distinguish remote proxy ARP from local proxy ARP, local proxy ARP performs a similar function but only when the requested IP is on the receiving interface.

Default **no remote-proxy-arp**

redundant-interface

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

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

Description This command configures a redundant interface used for dual homing.

Parameters *red-ip-int-name* — Specifies the redundant IP interface name.

arp-populate

Syntax [**no**] **arp-populate**

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

Description 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.

Enabling the **arp-populate** command will remove any dynamic ARP entries learned on this interface from the ARP cache.

The **arp-populate** command will fail if an existing static ARP entry exists for this interface.

The **arp-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

Syntax [**no**] **host ip** *ip-address* [**mac** *ieee-address*]] [**subscriber** *sub-ident-string*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**ancp-string** *ancp-string*]
no host {[**ip** *ip-address*] [**mac** *ieee-address*]}
no host all

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

Description This command creates a static subscriber host for the SAP. Static subscriber hosts may be used by the system for various purposes. Applications within the system that make use of static host entries include anti-spoof filters and ARP cache population.

Multiple static hosts may be defined on the SAP. Each host is identified by either a source IP address, a source MAC address or both a source IP and source MAC address. Every static host definition must have at least one address defined, IP or MAC.

Static hosts can exist on the SAP even with anti-spoof and ARP populate features disabled. When enabled, each feature has different requirements for static hosts.

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

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

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

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

Default **none**

Parameters **ip** *ip-address* — Specify this optional parameter when defining a static host. The IP address must be specified for **anti-spoof ip**, **anti-spoof ip-mac** and **arp-populate** Only one static host may be configured on the SAP with a given IP address.

mac *mac-address* — Specify this optional parameter when defining a static host. The MAC address must be specified for **anti-spoof ip-mac** and **arp-populate**. Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address.

subscriber *sub-ident-string* — Specify this optional parameter to specify an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

- For VPRN SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's sub-ident-string is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPRN destinations.

  If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

  If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.

  If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

  ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

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

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

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

IES Interface DHCP Commands

dhcp

Syntax **dhcp**

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

Description This command enables the context to configure DHCP parameters.

action

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

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

Description This command configures the Relay Agent Information Option (Option 82) processing.

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

Default The default is to keep the existing information intact.

Parameters **replace —** In the upstream direction (from the user), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.

keep — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is forwarded towards the client.

The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.

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

circuit-id

Syntax **circuit-id** [**ascii-tuple | ifindex | sap-id | vlan-ascii-tuple**]
no circuit-id

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

Description When enabled, the router sends either an ASCII tuple, or the interface index (If Index), on the specified SAP ID in the **circuit-id** suboption of the DHCP packet.

If disabled, the **circuit-id** suboption of the DHCP packet will be left empty.

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

Default **circuit-id ascii-tuple**

Parameters **ascii-tuple —** Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by "|".

ifindex — Specifies that the interface index will be used. The If Index of a router interface can be displayed using the command `show>router>interface>detail`.

sap-id — Specifies that the SAP ID will be used.

vlan-ascii-tuple — Specifies that the format will include VLAN ID, dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

match-circuit-id

Syntax [**no**] **match-circuit-id**

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

Description This command enables Option 82 circuit ID on relayed DHCP packet matching.

For Routed CO, the group interface DHCP relay process is stateful. When packets are relayed to the server the virtual router ID, transaction ID, SAP ID, and client hardware MAC address of the relayed packet are tracked. When a response is received from the server the virtual router ID, transaction ID, and client HW MAC address must be matched to determine the SAP on which to send the packet out. In some cases, the virtual router ID, transaction ID, and client HW MAC address are not guaranteed to be unique.

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

Default **no match-circuit-id**

option

Syntax [**no**] **option**

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

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

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

Default **no option**

remote-id

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

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

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

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

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

Default **remote-id**

Parameters **mac —** This keyword specifies the MAC address of the remote end is encoded in the suboption.

string *string* **—** Specifies the remote-id.

vendor-specific-option

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

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

Description This command configures the vendor specific suboption of the DHCP relay packet.

client-mac-address

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

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

Description This command enables the sending of the MAC address in the vendor specific suboption of the DHCP relay packet.

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

sap-id

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

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

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

Description This command enables the sending of the SAP ID in the vendor specific suboption of the DHCP relay packet.

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

service-id

Syntax [**no**] **service-id**

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

Description This command enables the sending of the service ID in the vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the vendor specific suboption of the DHCP relay packet.

string

Syntax [**no**] **string** *text*

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

Description This command specifies the string in the vendor specific suboption of the DHCP relay packet.

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

Parameters *text* — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

system-id

Syntax [**no**] **system-id**

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

Description This command specifies whether the system-id is encoded in the vendor specific sub-option of Option 82.

proxy-server

Syntax **proxy-server**

Context config>service>ies>if>dhcp

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

Description This command configures the DHCP proxy server.

emulated-server

Syntax **emulated-server** *ip-address*
no emulated-server

Context config>service>ies>if>dhcp>proxy-server
config>service>ies>sub-if>grp-if>dhcp>proxy-server

Description This command configures the IP address which will be used as the DHCP server address in the context of this SAP. Typically, the configured address should be in the context of the subnet represented by service.

The **no** form of this command reverts to the default setting. The local proxy server will not become operational without the emulated-server address being specified.

Parameters *ip-address* — Specifies the emulated server address.

lease-time

Syntax **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*] [**radius-override**]
no lease-time

Context config>service>ies>if>dhcp>proxy-server
config>service>ies>sub-if>grp-if>dhcp>proxy-server

Description This command defines the length of lease-time that will be provided to DHCP clients. By default the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.

The no form of this command disables the use of the lease-time command. The local-proxy-server will use the lease-time offered by either a RADIUS or DHCP server.

Default 7 days 0 hours 0 seconds

Parameters **radius-override —** Specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients.

days — Specifies the number of days that the given IP address is valid.

Values 0 — 3650

hours — Specifies the number of hours that the given IP address is valid.

Values 0 — 23

minutes — Specifies the number of minutes that the given IP address is valid.

Values 0 — 59

seconds — Specifies the number of seconds that the given IP address is valid.

Values 0 — 59

server

Syntax **server** *server1* [*server2*...(up to 8 max)]

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

Description This command specifies a list of servers where requests will be forwarded. The list of servers can be entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list.

There can be a maximum of 8 DHCP servers configured.

Default **no server**

Parameters *server* — Specify the DHCP server IP address.

trusted

Syntax [**no**] **trusted**

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

Description According to RFC 3046, *DHCP Relay Agent Information Option*, a DHCP request where the giaddr is 0.0.0.0 and which contains a Option 82 field in the packet, should be discarded, unless it arrives on a "trusted" circuit. If trusted mode is enabled on an IP interface, the Relay Agent (the router) will modify the request's giaddr to be equal to the ingress interface and forward the request.

Note that this behavior only applies when the action in the Relay Agent Information Option is "keep". In the case where the Option 82 field is being replaced by the Relay Agent (action = "replace"), the original Option 82 information is lost anyway, and there is thus no reason for enabling the trusted option.

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

Default **not enabled**

gi-address

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

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

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

By default, the GI address used in the relayed DHCP packet is the primary IP address of a normal IES interface. Specifying the GI address allows the user to choose a secondary address. For group

interfaces a GI address must be specified under the group interface DHCP context or subscriber-interface DHCP context in order for DHCP to function.

Default **no gi-address**

Parameters *ip-address* — Specifies the host IP address to be used for DHCP relay packets.

src-ip-address — Specifies that this GI address is to be the source IP address for DHCP relay packets.

lease-populate

Syntax **lease-populate** [*number-of-entries*]
no lease-populate

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

Description This command enables dynamic host lease state management for IES IP interfaces. Lease state information is extracted from snooped or relayed DHCP ACK messages to populate lease state table entries for the IP interface.

If lease state population is enabled and an entry cannot be added to the table, the system will prevent the far-end host from receiving the DHCP ACK message.

The retained lease state information representing dynamic hosts may be used to populate a SAP based anti-spoof filter table to provide dynamic anti-spoof filtering and/or ARP cache population.

Default **ies>if>dhcp: no lease-populate**

ies>sub-if>grp-if>dhcp: 1

Parameters *number-of-entries* — This optional parameter defines the number of lease state table entries allowed for this IP interface. If *number-of-entries* is not specified, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed.

IES Interface ICMP Commands

icmp

Syntax **icmp**

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

Description Context for configuring Internet Control Message Protocol (ICMP) parameters on an IES service

mask-reply

Syntax [**no**] **mask-reply**

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

Description This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance will reply to mask requests.

The **no** form of this command disables replies to ICMP mask requests on the router interface.

Default **mask-reply** - Reply to ICMP mask requests.

redirects

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

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

Description This commad configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.

When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.

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

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval. (*Default: redirects 100 10*)

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

unreachables

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

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

Description Enables/disables 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 *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 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

IES Interface IPv6 Commands

ipv6

Syntax [**no**] **ipv6**

Context config>service>ies>interface

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

address

Syntax **address** *ipv6-address/prefix-length* [**eui-64**]
no address *ipv6-address/prefix-length*

Context config>service>ies>if>ipv6

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

Parameters *ipv6-address/prefix-length* — Specify the IPv6 address on the interface.

Values		
	ipv6-address/prefix: ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x [0 — FFFF]H d [0 — 255]D
	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.

dhcp6-relay

Syntax [**no**] **dhcp6-relay**

Context config>service>ies>if>ipv6

Description This command enables the context to configure DHCPv6 relay parameters for the IES interface.

The **no** form of the command disables DHCP6 relay.

lease-populate

Syntax **lease-populate** [*nbr-of-entries*]
no lease-populate

Context config>service>ies>if>ipv6>dhcp-relay

Description This command specifies the maximum number of DHCP6 lease states allocated by the DHCP6 DHCP relay function, allowed on this interface.

The **no** form of the command disables dynamic host lease state management.

Default **no lease-populate**

Parameters *nbr-of-entries* — Defines the number lease state table entries allowed for this interface. If this parameter is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCP6 ACK messages are discarded.

Values 1 — 8000

neighbor-resolution

Syntax [**no**] **neighbor-resolution**

Context config>service>ies>if>ipv6>dhcp6-relay

Description This command enables neighbor resolution with DHCPv6 relay.

The **no** form of the command disables neighbor resolution.

option

Syntax [**no**] **option**

Context config>service>ies>if>ipv6>dhcp6-relay

Description This command enables the context to configure DHCPv6 relay information options.

The **no** form of the command disables DHCPv6 relay information options.

interface-id

Syntax **interface-id**
interface-id ascii-tuple
interface-id ifindex
interface-id sap-id
interface-id string
no interface-id

Context config>service>ies>if>ipv6>dhcp6>option

Description This command enables the sending of interface ID options in the DHCPv6 relay packet.

The **no** form of the command disables the sending of interface ID options in the DHCPv6 relay packet

Parameters **ascii-tuple —** Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by "|".

ifindex — Specifies that the interface index will be used. (The If Index of a router interface can be displayed using the command `show>router>interface>detail`)

sap-id — Specifies that the SAP identifier will be used.

string — Specifies a string of up to 32 characters long, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

remote-id

Syntax [**no**] **remote-id**

Context config>service>ies>if>ipv6>dhcp6>option

Description This command enables the sending of remote ID option in the DHCPv6 relay packet.

The client DHCP Unique Identifier (DUID) is used as the remote ID.

The **no** form of the command disables the sending of remote ID option in the DHCPv6 relay packet.

server

Syntax **server** *ipv6z-address* [*ipv6z-address*...(up to 8 max)]

Context config>service>ies>if>ipv6>dhcp6

Description This command specifies a list of servers where DHCP6 requests will be forwarded. The list of servers can entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP6 relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list.

There can be a maximum of 8 DHCP6 servers configured.

Default **no server**

Parameters *ipv6-address* — Specifies the IPv6 addresses of the DHCP servers where the DHCP6 requests will be forwarded. Up to 8 addresses can be specified.

Values ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x: [0 — FFFF]H
d: [0 — 255]D

source-address

Syntax **source-address** *ipv6-address*
no source-address

Context config>service>ies>if>ipv6>dhcp6

Description This command configures the source IPv6 address of the DHCPv6 relay messages.

Parameters *ipv6-address* — Specifies the source IPv6 address of the DHCPv6 relay messages.

Values ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x: [0 — FFFF]H
d: [0 — 255]D

dhcp6-server

dhcp6-server

Syntax [**no**] **dhcp6-server**

Context config>service>ies>if>ipv6

Description This command enables the context to configure DHCPv6 server parameters for the IES interface.

The **no** form of the command disables the DHCP6 server.

max-nbr-of-leases

Syntax **max-nbr-of-leases** *max-nbr-of-leases*
no max-nbr-of-leases

Context config>service>ies>if>ipv6>dhcp6-server

Description This command configures the maximum number of lease states installed by the DHCP6 server function allowed on this interface.

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

Default **8000**

Parameters *max-nbr-of-leases* — Specifies the maximum number of lease states installed by the DHCP6 server function allowed on this interface.

Values 0 — 8000

prefix-delegation

Syntax [**no**] **prefix-delegation**

Context config>service>ies>if>ipv6>dhcp6-server

Description This command configures prefix delegation options for delegating a long-lived prefix from a delegating router to a requesting router, where the delegating router does not require knowledge about the topology of the links in the network to which the prefixes will be assigned.

The **no** form of the command disables prefix-delegation.

prefix

Syntax [**no**] **prefix** *ipv6-address/prefix-length*

Context config>service>ies>if>ipv6>dhcp6-server>pfx-delegate

Description This command specifies the IPv6 prefix that will be delegated by this system.

Parameters *ipv6-address/prefix-length* — Specify the IPv6 address on the interface.

Values

ipv6-address/prefix: ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x [0 — FFFF]H d [0 — 255]D
prefix-length	1 — 128

duid

Syntax **duid** *duid* [**iaid** *iaid*]
no duid

Context config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix

Description This command configures the DHCP Unique Identifier (DUID) of the DHCP client.

Parameters *duid* — Specifies the ID of the requesting router. If set to a non zero value the prefix defined will only be delegated to this router. If set to zero, the prefix will be delegated to any requesting router.

iaid *iaid* **—** Specifies the identity association identification (IAID) from the requesting router that needs to match in order to delegate the prefix defined in this row.If set to 0 no match on the received IAID is done.

preferred-lifetime

Syntax **preferred-lifetime** *seconds*
preferred-lifetime infinite
no preferred-lifetime

Context config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix

Description This command configures the IPv6 prefix/mask preferred life time. The preferred-lifetime value cannot be bigger than the valid-lifetime value.

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

Default **604800 seconds** (7 days)

Parameters *seconds* — Specifies the time, in seconds, that this prefix remains preferred.

Values 1 — 4294967294

infinite — Specifies that this prefix remains preferred infinitely.

valid-lifetime

Syntax **valid-lifetime** *seconds*
valid-lifetime infinite
no valid-lifetime

Context config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix

Description This command configures the time, in seconds, that the prefix is valid. The maximum value 4294967295 is considered equal to infinity.

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

Default **2592000 seconds** (30 days)

Parameters *seconds* — Specifies the time, in seconds, that this prefix remains valid.

Values 1 — 4294967294

infinite — Specifies that this prefix remains valid infinitely.

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** [*number seconds*]
no packet-too-big

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

Description This command specifies whether “packet-too-big” ICMP messages should be sent. When enabled, ICMPv6 “packet-too-big” messages are generated by this interface.

The **no** form of the command disables the sending of ICMPv6 “packet-too-big” messages.

Default **100 10**

Parameters *number* — Specifies the number of “packet-too-big” ICMP messages to send 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 "packet-too-big" ICMP messages issued.

Values 1 — 60

Default 10

param-problem

Syntax **param-problem** [*number seconds*]
no packet-too-big

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

Description This command specifies whether "parameter-problem" ICMP messages should be sent. When enabled', "parameter-problem" ICMP messages are generated by this interface.

The **no** form of the command disables the sending of "parameter-problem" ICMP messages.

Default **100 10**

number — Specifies the number of "parameter-problem" ICMP messages to send 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 "parameter-problem" ICMP 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" ICMP 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" ICMP 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" ICMP message to be issued.

Values 1 — 60

Default 10

unreachables

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

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

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:x (eight 16-bit pieces)
x: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.

IES Spoke SDP Commands

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[:*vc-id*]
no spoke-sdp *sdp-id*[:*vc-id*]

Context config>service>ies>interface
config>service>ies>redundant-interface

Description This command binds a service to an existing Service Distribution Point (SDP).

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

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with an IES service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router. The spoke SDP must be shut down first before it can be deleted from the configuration.

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

Special Cases **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

egress

Syntax **egress**

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

Description This command configures the egress SDP context.

vc-label

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

Context config>service>ies>if>spoke-sdp>egress
config>service>ies>redundant-interface>spoke-sdp>egress

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

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

Values 16 — 1048575

ingress

Syntax **ingress**

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

Description This command configures the ingress SDP context.

vc-label

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

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

IES SAP Commands

sap

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

Context config>service>ies>interface
config>service>ies>sub-if>grp-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 7750. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface** *port-type port-id* **mode access** command. Channelized TDM ports are always access ports.

If a port is 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 Ethernet Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed.

Default No SAPs are defined.

Special Cases **IES —** An IES SAP can be defined with Ethernet ports, SONET/SDH or TDM channels. A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. Group interfaces allow more than one SAP. 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.

The *sap-id* can be configured in one of the following formats:

Type	Syntax	Example
port-id	*slot*/*mda*/*port*[.*channel*]	2/2/11 1/2/3.1
null	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id* | *aps-id*]	*port-id*: 1/2/3 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*: lag-100 *aps-id*: aps-1
dot1q	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id* | *aps-id*]:qtag1	*port-id*:qtag1: 1/2/3:100 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*:qtag1:lag-100:102 *aps-id*:qtag1: aps-1:103
qinq	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	*port-id*:qtag1.qtag2: 1/2/3:100.10 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*:qtag1.qtag2:lag-100:
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:vpi/vci |vpi |vpi1.vpi2]	port-id: 9/1/1 aps-id: aps-1 *bundle-id*: bundle-ima-1/1.1 bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 vpi/vci: 16/26 vpi: 16 vpi1.vpi2: 16.200
frame-relay	[*port-id* | *aps-id*]:*dlci*	*port-id*: 1/1/1:100 *aps-id*: aps-1 *dlci*: 16
cisco-hdlc	*slot/mda/port.channel*	*port-id*: 1/2/3.1

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	

	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

bundle-id — Specifies the multilink bundle to be associated with this IP interface. The **bundle** keyword must be entered at the beginning of the parameter.
The command syntax must be configured as follows:

bundle-id: **bundle-***type-slot-id*/*mda-slot*.*bundle-num*
bundle-id value range: 1 — 128

For example:

```
*A:ALA-12>config#  port bundle-ppp-5/1.1
*A:ALA-12>config>port# multilink-bundle
```

bgprp-id — Specifies the bundle protection group ID to be associated with this IP interface. The **bpgrp** keyword must be entered at the beginning of the parameter.

The command syntax must be configured as follows:

bpgrp-id: **bpgrp**-*type-bpgrp-num*
type: ima
bpgrp-num value range: 1 — 1280

For example:

```
*A:ALA-12>config# port bpgrp-ima-1
*A:ALA-12>config>service>vpls$ sap bpgrp-ima-1
```

qtag1, qtag2 — Specifies the encapsulation value used to identify the SAP on the port or sub-port. If this parameter is not specificially defined, the default value is 0.

Values

qtag1:	0 — 4094
qtag2 :	* | 0 — 4094

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 TDM	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 TDM	BCP-Dot1q	0 — 4094	The SAP is identified by the 802.1Q tag on the channel.
SONET/SDH TDM	Frame Relay	16 — 991	The SAP is identified by the data link connection identifier (DLCI).
SONET/SDH ATM	ATM	vpi (NNI) 0 — 4095 vpi (UNI) 0 — 255 vci 1, 2, 5 — 65535 -	The SAP is identified by port or by PVPC or PVCC identifier (vpi, vpi/vci, or vpi range)

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

multi-service-site

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

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

Description This command creates a new customer site or edits an existing customer site with the *customer-site-name* parameter. A customer site is an anchor point to create an ingress and egress virtual scheduler hierarchy. When a site is created, it must be assigned to a chassis slot or port with the exception of the

7750 SR-1 in which the slot is set to 1. When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).

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.

tod-suite

Syntax **tod-suite** *tod-suite-name*
no tod-suite

Context config>service>ies>if>sap
config>service>ies>sub-if>grp-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>cron** context.

Default **no tod-suite**

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

host-connectivity-verify

Syntax **host-connectivity-verify** [**source {vrrp | interface**}] [**interval** *interval*] [**action** {**remove | alarm**}]

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

Description This command enables subscriber host connectivity verification for all hosts on this interface. This tool will periodically scan all known hosts (from dhcp-state) and perform a UC ARP request. The subscriber host connectivity verification will maintain state (connected vs. not-connected) for all hosts.

Default **no host-connectivity-verify**

Parameters **source** {**interface**} **—** Specifies the source to be used for generation of subscriber host connectivity verification packets. The **interface** keyword forces the use of the interface mac and ip addresses. 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 *interval* — The interval, in minutes, which specifies the time interval which all known sources should be verified. The actual rate is then dependent on number of known hosts and interval.

Values 1 — 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.

SAP Subscriber Management Commands

sub-sla-mgmt

Syntax [**no**] **sub-sla-mgmt**

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

Description This command enables the context to configure subscriber management parameters for this SAP.

Default **no sub-sla-mgmt**

def-sla-profile

Syntax **def-sla-profile** *default-sla-profile-name*
no def-sla-profile

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

Description This command specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sla-profile** context.

An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts. SLA profiles are maintained in two locations, the subscriber identification policy and the subscriber profile templates. After a subscriber host is associated with an SLA profile name, either the subscriber identification policy used to identify the subscriber or the subscriber profile associated with the subscriber host must contain an SLA profile with that name. If both the subscriber identification policy and the subscriber profile contain the SLA profile name, the SLA profile in the subscriber profile is used.

The **no** form of the command removes the default SLA profile from the SAP configuration.

Default **no def-sla-profile**

Parameters *default-sla-profile-name* — Specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sla-profile** context.

def-sub-profile

Syntax **def-sub-profile** *default-subscriber-profile-name*

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

Description This command specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-profile** context.

A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for

subscriber using the subscriber profile. Subscriber profiles also allow for specific SLA profile definitions when the default definitions from the subscriber identification policy must be overridden.

The **no** form of the command removes the default SLA profile from the SAP configuration.

Parameters *default-sub-profile* — Specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-profile** context.

sub-ident-policy

Syntax **sub-ident-policy** *sub-ident-policy-name*

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

Description This command associates a subscriber identification policy to this SAP. The subscriber identification policy must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-ident-policy** context.

Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.

For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet string. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.

When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.

The **no** form of the command removes the default subscriber identifcation policy from the SAP configuration.

Default **no sub-ident-policy**

Parameters *sub-ident-policy-name* — Specifies a subscriber identification policy for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-ident-policy** context.

multi-sub-sap

Syntax **multi-sub-sap** [*subscriber-limit*]
no multi-sub-sap

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

Description This command configures the maximum number of subscribers for this SAP.

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

Parameters *subscriber-limit* — Specifies the maximum number of subscribers for this SAP.

Values 2 — 8000

single-sub-parameters

Syntax **single-sub-parameters**

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

Description This command enables the context to configure single subscriber parameters for this SAP.

non-sub-traffic

Syntax **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*]
no non-sub-traffic

Context config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt>single-sub

Description This command configures non-subscriber traffic profiles.

The **no** form of the command removes removes the profiles and disables the feature.

Parameters **sub-profile** *sub-profile-name* **—** Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

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

subscriber *sub-ident-string* — Specifies an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

- For VPRN SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's sub-ident-string is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPRN destinations.

  If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

  If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.

  If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

  ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

profiled-traffic-only

Syntax [**no**] **profiled-traffic-only**

Context config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt>single-sub

Description This command enables profiled traffic only for this SAP.

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

Service Billing Commands

accounting-policy

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

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

Description This command creates the accounting policy context that can be applied to a SAP.

An accounting policy must be defined before it can be associated with a SAP.
If the *policy-id* does not exist, an error message is generated.

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

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

Default Default accounting policy.

Parameters *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config**>**log**>**accounting-policy** context.

Values 1 to 99

collect-stats

Syntax [**no**] **collect-stats**

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

Description This command enables accounting and statistical data collection for either the SAP, network port, or IP interface. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default **collect-stats**

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
config>service>ies>redundant-interface>egress
config>service>ies>redundant-interface>ingress
config>service>ies>sub-if>grp-if>sap>egress
config>service>ies>sub-if>grp-if>sap>ingress

Description This command associates afilter 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* or *ipv6-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**.

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

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

ip-filter-id — Specifies the ID for the IP filter policy. Allowed values are an integer 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
no filter [**ip** *ip-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**.

Special Cases **IES —** Only IP filters are supported on IES IP interfaces, and the filters only apply to routed traffic.

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

ip-filter-id — The filter name acts as the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy. The filter ID must already exist within the created IP filters.

egress

Syntax **egress**

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

Description This command enables the context to configure egress SAP Quality of Service (QoS) policies and filter policies.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing. If no egress filter is defined, no filtering is performed.

ingress

Syntax **ingress**

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

Description This command enables the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

match-qinq-dot1p

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

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

Description This command specifies which Dot1Q tag position Dot1P bits in a QinQ encapsulated packet should be used to evaluate Dot1P QoS classification.

The **match-qinq-dot1p** command allows the top or bottom PBits to be used when evaluating the applied sap-ingress QoS policy's Dot1P entries. The **top** and **bottom** keywords specify which position should be evaluated for QinQ encapsulated packets.

The **no** form of the command restores the default dot1p evaluation behavior for the SAP.

By default, the bottom most service delineating Dot1Q tags Dot1P bits are used. Table 24 defines the default behavior for Dot1P evaluation when the **match-qinq-dot1p** command is not executed.

Table 1: 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 25 defines the dot1p evaluation behavior when the top parameter is specified.

Table 2: 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 26 defines the dot1p evaluation behavior when the bottom parameter is specified.

Table 3: 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 3: 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 4: 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 5: 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

Table 5: QinQ Mark Top Only Explicit Marking Actions

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
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 PBits marked using dot1p-value, BottomQ PBits marked with zero
QinQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits PBits marked using dot1p-value, BottomQ PBits marked using preserved value

The QinQ and TopQ SAP PBit marking follows the default behavior devined in Table 27 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*
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, at 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

qinq-mark-top-only

Syntax [**no**] **qinq-mark-top-only**

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

Description When enabled (the encapsulation type of the access port where this SAP is defined as qinq), the **qinq-mark-top-only** command specifies which P-bits to mark during packet egress. When disabled, both set of P-bits are marked. When the enabled, only the P-bits in the top Q-tag are marked.

Default **no qinq-mark-top-only**

qos

Syntax **qos** *policy-id* [**shared-queuing** |**multipoint-shared**]
no qos

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

Description Associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP) or IP interface.

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the *policy-id* does not exist, an error will be returned.

The **qos** command is used to associate both ingress and egress QoS policies. The **qos** command only allows ingress policies to be associated on SAP or IP interface ingress and egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a 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.

When an egress QoS policy is associated with an IES IP interface that has been bound to a VPLS, the policy becomes associated with every SAP on the VPLS and augments the egress 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 or IP interface for ingress or egress, so the default QoS policy is used.

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

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 — This keyword can only be specified on SAP ingress. Specify 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. If you attempt to map forwarding class unicast traffic to a multipoint queue, an error is generated and 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.

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 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** *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.

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.

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.

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.

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.

avg-frame-overhead

Syntax **avg-frame-overhead** *percent*
no avg-frame-overhead

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

Description This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a Sonet or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- Offered-load — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- Frame encapsulation overhead — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be 10000 x 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 x 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 x 1.1 or 550 octets.
- Frame based within-cir offered-load — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

  As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.
- Frame based PIR — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be 7500 x 1.1 or 8250 octets.
- Frame based within-pir offered-load — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

Port scheduler operation using frame transformed rates — The port scheduler uses the frame based rates to determine the maximum rates that each queue may receive during the within-cir and above-cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

SAP and subscriber SLA-profile average frame overhead override — The average frame overhead parameter on a sap-egress may be overridden at an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers an avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

Default **0**

Parameters *percent* — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

Values 0 — 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 or 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>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 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

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

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

Description This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled 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 — 100000000

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.

Values 0 — 100000000, **max**, **sum**

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

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 then 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 ate 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. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues.

Values 0 — 10,000,000, **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
config>service>ies>sub-if>grp-if>sap>egress
config>service>ies>sub-if>grp-if>sap>ingress
config>service>ies>sub-if>grp-if>sap>egress
config>service>ies>sub-if>grp-if>sap>ingress

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.

ATM Commands

atm

Syntax **atm**

Context config>service>ies>if>sap

Description This command enables access to the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as:

- Configuring ATM port or ATM port-related functionality on MDAs supporting ATM functionality
- Configuring ATM-related configuration for ATM-based SAPs that exist on MDAs supporting ATM functionality.

If ATM functionality is not supported for a given context, the command returns an error.

egress

egress

Context config>service>ies>if>sap>atm

Description This command enables the context to configure egress ATM attributes for the SAP.

encapsulation

Syntax **encapsulation** *atm-encap-type*

Context config>service>ies>if>sap>atm

Description This command configures RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, encapsulation for an ATM PVCC delimited SAP.

This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition references RFC 2684 and to the ATM Forum LAN Emulation specification.

Ingress traffic that does not match the configured encapsulation will be dropped.

Default The encapsulation is driven by the services for which the SAP is configured.
For IES service SAPs, the default is **aal5snap-routed**.

Parameters *atm-encap-type* — Specify the encapsulation type.

Values **aal5snap-routed** — Routed encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684.
aal5mux-ip — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684.

aal5snap-bridged — Bridged encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684.
aal5mux-bridged-eth-nofcs — Bridged IP encapsulation for VC multiplexed circuit as defined in RFC 2684.

ingress

Syntax **ingress**

Context config>service>ies>if>sap>atm

Description This command configures ingress ATM attributes for the SAP.

traffic-desc

Syntax **traffic-desc** *traffic-desc-profile-id*
no traffic-desc

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

Description This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP).

When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction.

When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.

The **no** form of the command reverts the traffic descriptor to the default traffic descriptor profile.

Default The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.

Parameters *traffic-desc-profile-id* — Specify a defined traffic descriptor profile (see the QoS atm-td-profile command).

oam

Syntax **oam**

Context config>service>ies>if >sap>atm

Description This command enables the context to configure OAM functionality for a PVCC delimiting a SAP.

The ATM-capable MDAs support F5 end-to-end OAM functionality (AIS, RDI, Loopback):

- ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95
- GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996

- GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

alarm-cells

Syntax [**no**] **alarm-cells**

Context config>service>ies>if >sap>atm>oam

Description This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC termination to monitor and report the status of their connection by propagating fault information through the network and by driving PVCC's operational status.

When alarm-cells functionality is enabled, a PVCC's operational status is affected when a PVCC goes into an AIS or RDI state because of an AIS/RDI processing (assuming nothing else affects PVCC's operational status, for example, if the PVCC goes DOWN, or enters a fault state and comes back UP, or exits that fault state). RDI cells are generated when PVCC is operationally DOWN. No OAM-specific SNMP trap is raised whenever an endpoint enters/exits an AIS or RDI state, however, if as result of an OAM state change, the PVCC changes operational status, then a trap is expected from an entity the PVCC is associated with (for example a SAP).

The **no** command disables alarm-cells functionality for a PVCC. When alarm-cells functionality is disabled, a PVCC's operational status is no longer affected by a PVCC's OAM state changes due to AIS/RDI processing (note that when alarm-cells is disabled, a PVCC will change operational status to UP due to alarm-cell processing) and RDI cells are not generated as result of the PVCC going into AIS or RDI state. The PVCC's OAM status, however, will record OAM faults as described above.

Default Enabled for PVCCs delimiting IES SAPs

periodic-loopback

Syntax [**no**] **periodic-loopback**

Context config>service>ies>if >sap>atm>oam

Description This command enables periodic OAM loopbacks on this SAP. This command is only configurable on IES and VPRN SAPs. When enabled, an ATM OAM loopback cell is transmitted every period as configured in the `config>system>atm>oam>loopback-period` *`period`* context.

If a response is not received and consecutive retry-down retries also result in failure, the endpoint will transition to an alarm indication signal/loss of clock state. Then, an ATM OAM loopback cell will be transmitted every period as configured in the `loopback-period` *`period`*. If a response is received for the periodic loopback and consecutive retry-up retries also each receive a response, the endpoint will transition back to the up state.

The **no** form of the command sets the value back to the default.

Default **no periodic-loopback**

IES Interface VRRP Commands

vrrp

Syntax **vrrp** *virtual-router-id* [**owner**]
no vrrp *virtual-router-id*

Context config>service>ies>interface

Description This command creates or edits a Virtual Router ID (VRID) on the service IP interface. A VRID is internally represented in conjunction with the IP interface name. This allows the VRID to be used on multiple IP interfaces while representing different virtual router instances.

Two VRRP nodes can be defined on an IP interface. One, both, or none may be defined as owner. The nodal context of **vrrp** *virtual-router-id* is used to define the configuration parameters for the VRID.

The **no** form of this command removes the specified VRID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the vrid. The VRID does not need to be shutdown in order to remove the virtual router instance.

Default **No default**

Parameters *virtual-router-id* — The virtual-router-id parameter specifies a new virtual router ID or one that can be modified on the IP interface.

Values 1 — 255

authentication-key

Syntax **authentication-key** [*authentication-key* **|** *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>service>ies>if>vrrp

Description The **authentication-key** command, within the **vrrp** *virtual-router-id* context, is used to assign a simple text password authentication key to generate master VRRP advertisement messages and validating received VRRP advertisement messages.

The authentication-key command is one of the few commands not affected by the presence of the owner keyword. If simple text password authentication is not required, the authenticaton-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, altering the simple text password used when authentication-type password authentication method is used by the virtual router instance. The authentication-type password command does not need to be executed prior to defining the authentication-key command.

To change the current in-use password key on multiple virtual router instances:

- Identify the current master
- Shutdown the virtual router instance on all backups

- Execute the authentication-key command on the master to change the password key
- Execute the authentication-key command and no shutdown command on each backup key

The **no** form of this command restores the default null string to the value of key.

Default No default. The authentication data field contains the value 0 in all 16 octets.

Parameters *authentication-key* — 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.

authentication-type

Syntax **authentication-type** {*password* | *message-digest*}
no authentication-type

Context config>service>ies>if>vrrp

Description The **authentication-type** command, within the **vrrp** *virtual-router-id* context, is used to assign the authentication method to generate master VRRP advertisement messages and validate received VRRP advertisement messages.

NOTE: The authentication management for VRRP closely follows the authentication management format used for IS-IS.

The **authentication-type** command is one of the commands not affected by the presence of the owner keyword. If authentication is not required, the authenticaton-type command must not be executed. If the command is re-executed with a different authentication type defined, the new type will be used. If the no authentication-type command is executed, authentication is removed and no authentication is performed. The authentication-type command may be executed at any time, altering the authentication method used by the virtual router instance.

The **no** form of this command removes authentication from the virtual router instance. All VRRP Advertisement messages sent will have the Authentication Type field set to 0 and the Authentication Data fields will contain 0 in all octets. VRRP Advertisement messages received with Authentication Type fields containing a value other than 0 will be discarded.

password — The password keyword identifies VRRP Authentication Type 1. Type 1 requires the definition of a string of eight octets long using the authentication-key command. All transmitted VRRP Advertisement messages must have the Authentication Type field set to 1 and the Authentication Data fields must contain the authentication-key password.

All received VRRP advertisement messages must contain a value of 1 in the Authentication Type field and the Authentication Data fields must match the defined authentication-key. All other received messages will be silently discarded.

message-digest — The message-digest keyword identifies VRRP Authentication Type 2. Type 2 defines a lower IP layer MD5 authentication mechanism using HMAC and IP authentication header standards. An MD5 key must be defined using the message-digest-key command. All transmitted VRRP advertisement messages must have the Authentication Type field set to 2 and the Authentication Data fields must contain 0 in all octets. The message-digest key is used in the hashing process when populating the IP Authentication Header fields. A sequential incrementing counter (set to zero when the message-digest-key is set) is incremented and then used in the IP Authentication Header to prevent replay attacks on authorized participating virtual router instances.

All received VRRP advertisement messages must contain a value of 2 in the Authentication Type field and the Authentication Data fields are ignored. The message must have been authorized by the lower layer IP Authentication Header process with the sequential counter field and the source IP address presented to the virtual router instance. To track the validity of the received counter, the virtual router instance maintains a master counter table containing up to 32 source IP addresses and the last received counter value. Populate the table as follows:

1. Check to see if source IP address exists in table.
   - If non-existent, create an entry if available.
   - If no entry is available, delete the oldest and create an entry.
     The new entry should have a counter value of zero.
2. Compare the message counter value to the entry value (0 if new entry or equal to the previous message counter from the source IP address).
   - If the message counter is not greater than the entry counter value, silently discard the packet.
   - If the message counter is greater than the entry counter value, accept the message for further checking and replace the entry counter value with the message counter value and time stamp the entry.

backup

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

Context config>service>ies>if>vrrp

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

init-delay

Syntax **init-delay** *seconds*
no init-delay

Context config>service>ies>if>vrrp

Description This command configures a VRRP initialization delay timer.

Default **no init-delay**

Parameters *seconds* — Specifies the initialization delay timer for VRRP, in seconds.

Values 1 — 65535

mac

Syntax **mac** *ieee-address*
no mac

Context config>service>ies>if>vrrp

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

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

Default The physical MAC address associated with the Ethernet interface that the SAP is configured on (the default MAC address assigned to the interface, assigned by the system).

Parameters *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.

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 This command sets the advertisement timer and indirectly sets the master down timer on the virtual router instance. The message-interval setting must be the same for all virtual routers participating as a virtual router. Any VRRP advertisement message received with an Advertisement Interval field different than the virtual router instance configured message-interval value will be silently discarded.

The message-interval command is available in both non-owner and owner **vrrp** *virtual-router-id* nodal contexts. If the message-interval command is not executed, the default message interval of 1 second will be used.

The **no** form of this command restores the default message interval value of 1 second to the virtual router instance.

Parameters *seconds* — The number of seconds that will transpire before the advertisement timer expires.

Values 1 — 255

Default 1

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

Values 100 — 900

ping-reply

Syntax **ping-reply**
no ping-reply

Context config>service>ies>if>vrrp

Description This command enables the non-owner master to reply to ICMP Echo Requests directed at the virtual router instances IP addresses. The ping request can be received on any routed interface.

Ping must not have been disabled at the management security level (either on the parental IP interface or based on the Ping source host address). When ping-reply is not enabled, ICMP Echo Requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to ICMP Echo Requests regardless of the setting of ping-reply configuration.

The ping-reply command is only available in non-owner **vrrp** *virtual-router-id* nodal context. If the ping-reply command is not executed, ICMP Echo Requests to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all ICMP Echo Request messages destined to the non-owner virtual router instance IP addresses.

Default no ping-reply

policy

Syntax **policy** *vrrp-policy-id*
no policy

Context config>service>ies>if>vrrp

Description This command creates VRRP control policies. The VRRP policy ID must be created by the policy command prior to association with the virtual router instance.

The policy command provides the ability to associate a VRRP priority control policy to a virtual router instance. The policy may be associated with more than one virtual router instance. The priority events within the policy either override or diminish the base-priority dynamically affecting the in-use priority. As priority events clear in the policy, the in-use priority may eventually be restored to the base-priority value.

The policy command is only available in the non-owner **vrrp** *virtual-router-id* nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed by VRRP priority control policies. For non-owner virtual router instances, if the policy command is not executed, the base-priority will be used as the in-use priority.

The **no** form of this command removes any existing VRRP priority control policy association from the virtual router instance. All such associations must be removed prior to the policy being deleted from the system.

Default 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.

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 **ssh-reply**

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 **telnet-reply**
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 **traceroute-reply** status.

Default **no traceroute-reply**

Show Commands

customer

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

Context show>service

Description 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 : VPLS Customer
Phone       : (567) 555-1212

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

Customer-ID : 274
Contact     : Mssrs. Beaucoup
Description : ABC Company
Phone       : 650 123-4567

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

-------------------------------------------------------
Total Customers : 8
-------------------------------------------------------
*A: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#
```

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.

Table 6: 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
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#
```

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 2049 - 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
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#
```

sap-using

Syntax **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 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.

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-name* **—** The session authentication policy for which to display matching SAPs.

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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 service identifier.
SapMTU	The SAP MTU value.
I.QoS	The SAP ingress QoS policy number specified on the ingress SAP.

I.MAC/IP	The MAC or IP filter policy ID applied to the ingress SAP.
E.QoS	The SAP egress QoS policy number specified on the egress SAP.
E.Mac/IP	The MAC or IP filter policy ID applied to the egress SAP
A.Pol	The accounting policy ID assigned to the SAP.
Adm	The desired state of the SAP.
Opr	The actual state of the SAP.

Sample Output

```
*A:ALA-12# show service sap-using sap 1/1
===============================================================================
Service Access Points
===============================================================================
PortId          SvcId      SapMTU I.QoS I.Mac/IP E.QoS E.Mac/IP A.Pol  Adm  Opr
-------------------------------------------------------------------------------
1/1/7:0         1          1518   10    8        10    none     none   Up   Up
1/1/11:0        100        1514   1     none     1     none     none   Down Down
1/1/7:300       300        1518   10    none     10    none     1000   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 3
-------------------------------------------------------------------------------
*A:ALA-12#

*A:ALA-12#  show service sap-using egress atm-td-profile 2
===============================================================================
Service Access Point Using ATM Traffic Profile 2
===============================================================================
PortId     SvcId    I.QoS I.Fltr E.QoS E.Fltr A.Pol Adm Opr
-------------------------------------------------------------------------------
5/1/1:0/11 511111    2    none     2    none    none  Up   Up
5/1/1:0/12 511112    2    none     2    none    none  Up   Up
5/1/1:0/13 511113    2    none     2    none    none  Up   Up
5/1/1:0/14 511114    2    none     2    none    none  Up   Up
5/1/1:0/15 511115    2    none     2    none    none  Up   Up
5/1/1:0/16 511116    2    none     2    none    none  Up   Up
5/1/1:0/17 511117    2    none     2    none    none  Up   Up
5/1/1:0/18 511118    2    none     2    none    none  Up   Up
5/1/1:0/19 511119    2    none     2    none    none  Up   Up
5/1/1:0/20 511120    2    none     2    none    none  Up   Up
5/1/1:0/21 511121    2    none     2    none    none  Up   Up
5/1/1:0/22 511122    2    none     2    none    none  Up   Up
5/1/1:0/23 511123    2    none     2    none    none  Up   Up
5/1/1:0/24 511124    2    none     2    none    none  Up   Up
5/1/1:0/25 511125    2    none     2    none    none  Up   Up
...
===============================================================================
*A:ALA-12#
```

sdp

Syntax **sdp** [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]

Context show>service

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 desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Opr MTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
IP address	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Adm Admin State	Specifies the desired state of the SDP.
Opr Oper State	Specifies the operating state of the SDP.
Deliver	Specifies the type of delivery used by the SDP: GRE or MPLS.
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.

Label	Description
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SDP.
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified.
Hello Time	Specifies how often the SDP Echo Request messages are transmitted on this SDP.
Deliver/Delivered	Specifies the type of delivery used by the SDP: GRE or MPLS.
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 amount of time to wait before the keepalive operating status is eligible to enter the alive state.
TX Hello Msgs	Specifies the number of SDP echo request messages transmitted since the keepalive was administratively enabled or the counter was cleared.
Rx Hello Msgs	Specifies the number of SDP echo request messages received since the keepalive was administratively enabled or the counter was cleared.
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
===============================================================================
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:ALA-12# show service sdp 8
===============================================================================
Service Destination Point (Sdp Id : 8)
===============================================================================
SdpId    Adm MTU   Opr MTU   IP address       Adm  Opr         Deliver Signal
-------------------------------------------------------------------------------
8        4462      4462      10.10.10.104     Up   Dn NotReady MPLS    TLDP
===============================================================================
*A:ALA-12#
===============================================================================
Service Destination Point (Sdp Id : 8) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 8  -(10.10.10.104)
-------------------------------------------------------------------------------
Description           : MPLS-10.10.10.104
SDP Id               : 8
Admin Path MTU       : 0                    Oper Path MTU      : 0
Far End              : 10.10.10.104         Delivery           : MPLS
Admin State          : Up                   Oper State         : Down
Flags                : SignalingSessDown TransportTunnDown
Signaling            : TLDP                 VLAN VC Etype      : 0x8100
Last Status Change   : 02/01/2007 09:11:39  Adv. MTU Over.     : No
Last Mgmt Change     : 02/01/2007 09:11:46
```

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

Associated LSP LIST :
Lsp Name             : to-104
Admin State          : Up                    Oper State         : Down
Time Since Last Tran*: 01d07h36m
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:ALA-12#
```

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 .
I.Label	The label used by the far-end device to send packets to this device in this service by this SDP.
E.Label	The label used by this device to send packets to the far-end device in this service by this SDP.

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#
```

service-using

Syntax **service-using** [**ies**] [**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 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-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#
```

subscriber-using

Syntax **subscriber-using** [**service-id** *service-id*] [**sap-id** *sap-id*] [**interface** *ip-int-name*] [**ip** *ip-address[/mask]*] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*]

Context show>service>subscriber-using

Description Displays subscribers using certain options.

Parameters **service-id *service-id* —** Display subscriber information about the specified service ID.

Values 1 — 2147483647

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port*.*channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type*-*slot/mda*.*bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type*-*bpgrp-num*	
	bpgrp	keyword

	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

interface *ip-int-name* — Display subscriber information about the specified interface.

ip *ip-address[/mask* — Display subscriber information about the specified IP address.

mac *ieee-address* — Display subscriber information about the specified MAC address.

sub-profile *sub-profile-name* — Display subscriber information about the specified subscriber profile name.

sla-profile *sla-profile-name* — Display subscriber information about the specified SLA profile name.

id

Syntax **id** *service-id* {**all** | **arp** | **base** | **sap** | **sdp**}

Context show>service

Description Display 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.

interface — Display service interfaces.

sap — Display SAPs associated to the service.

sdp — Display SDPs associated with the service.

all

Syntax **all**

Context show>service>id

Description 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.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.
Type	Indicates whether this Service SDP binding is a spoke or a mesh.
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.

Label	Description (Continued)
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the keepalive protocol.
Oper State	The current status of the keepalive protocol.
Hello Time	Specifies how often the SDP Echo Request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP Echo Request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
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 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.

Label	Description (Continued)
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 of.
Ingress sched-policy	Indicates the ingress QoS scheduler for the SAP.
Egress sched-policy	Indicates the egress QoS scheduler for the SAP.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
SAP Statistics	
Dropped	The number of packets or octets dropped.
Offered Hi Priority	The number of high priority packets, as determined by the SAP ingress QoS policy.
Offered Low Priority	The number of low priority packets, as determined by the SAP ingress QoS policy.
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
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.

Label	Description (Continued)
High priority offered	The packets or octets count of the high priority traffic for the SAP.
High priority dropped	The number of high priority traffic packets/octets dropped.
Low priority offered	The packets or octets count of the low priority traffic.
Low priority dropped	The number of low priority traffic packets/octets dropped.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
Out profile forwarded	The number of out-of-profile octets (rate above CIR) forwarded.
Egress Queue 1	The index of the egress QoS queue of the SAP.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
In profile dropped	The number of in-profile packets or octets dropped for the SAP.
Out profile forwarded	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Out profile dropped	The number of out-of-profile packets or octets discarded.

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 for a subscriber interface are displayed uniquely. Each MAC associated with the subscriber interface child group-interfaces are displayed with each subscriber interface ARP entry. They do not reflect actual ARP entries but are displayed along the subscriber 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.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

port-id — **interface —** 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.

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

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.
Expiry	The age of the ARP entry.
Interface	The interface applied to the service.
SAP	The SAP id.

Sample Output

```
A:ALA-49# show service id 88 arp
===============================================================================
ARP Table
===============================================================================
IP Address   MAC Address        Type    Expiry       Interface   SAP
-------------------------------------------------------------------------------
11.30.1.1    76:1e:ff:00:01:41  Other   00h00m00s    ies30       lag-1:30
11.31.1.1    76:1e:ff:00:01:41  Other   00h00m00s    ies30       lag-1:30
11.37.1.1    00:00:00:00:00:00  Other   00h00m00s    foo2        n/a
11.20.1.1    76:1e:ff:00:00:00  Other   00h00m00s    s2          subscrib*
             76:1e:ff:00:01:41                       g3          lag-1
11.20.1.10   00:00:aa:aa:aa:dd  Managed 00h00m00s    g3          lag-1:11
11.20.1.11   00:00:aa:aa:aa:dd  Managed 00h00m00s    g3          lag-1:11
11.20.1.12   00:00:aa:aa:aa:dd  Managed 00h00m00s    g3          lag-1:11
11.38.1.1    76:1e:ff:00:00:00  Other   00h00m00s    s3          subscrib*
             76:21:04:01:00:01                       g5          4/1/1
             76:21:04:01:00:01                       g7          4/1/1
11.39.1.1    76:1e:ff:00:00:00  Other   00h00m00s    s3          subscrib*
             76:21:04:01:00:01                       g5          4/1/1
             76:21:04:01:00:01                       g7          4/1/1
11.38.1.2    76:22:07:01:00:01  Managed 00h00m00s    g7          4/1/1:25*
11.38.10.1   76:22:07:01:00:01  Managed 00h00m00s    g7          4/1/1:25*
11.38.99.1   76:22:07:01:00:01  Managed 00h00m00s    g7          4/1/1:25*

===============================================================================
* indicats that the corresponding row element may have been truncated.
A:ALA-49#
```

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 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.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

Sample Output

```
*A:ALA-1# show service id 11 authentication statistics
===============================================================
Authentication statistics
===============================================================
```

```
Interface / SAP                Authentication  Authentication
                               Successful      Failed
---------------------------------------------------------------
vpls-11-90.1.0.254              1582            3
---------------------------------------------------------------
Number of entries: 1
===============================================================
*A:ALA-1#
```

base

Syntax **base**

Context show>service>id

Description 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#
```

dhcp

Syntax **dhcp**

Context show>service>id

Description This command enables the context to display DHCP information for the specified service.

lease-state

Syntax **lease-state** [[**sap** *sap-id*] | [**sdp** *sdp-id:vc-id*] | [**interface** *interface-name*] | [**ip-address** *ip-address*]] [**detail**]

Context show>service>id>dhcp

Description This command displays DHCP lease state related information.

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp-id — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Values 1 — 4294967295

interface *interface-name* — Displays information for the specified IP interface.

ip-address *ip-address* — Displays information associated with the specified IP address.

detail — Displays detailed information.

Sample Output

```
A:ALA-_Dut-A# show service id 13 dhcp lease-state
===============================================================================
DHCP lease state table, service 13
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
13.13.40.1      00:00:00:00:00:13 1/1/1:13            00h00m58s  Radius
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-_Dut-A#


A:ALA-_Dut-A# show service id 13 dhcp lease-state detail
===============================================================================
DHCP lease states for service 13
===============================================================================
Service ID           : 13
IP Address           : 13.13.40.1
Mac Address          : 00:00:00:00:00:13
Interface            : ies-13-13.13.1.1
SAP                  : 1/1/1:13
Remaining Lifetime   : 00h00m58s
Persistence Key      : N/A

Sub-Ident            : "Belgacom"
Sub-Profile-String   : "ADSL GO"
SLA-Profile-String   : "BE-Video"
Lease ANCP-String    : ""

Sub-Ident origin     : Radius
Strings origin       : Radius
Lease Info origin    : Radius

Ip-Netmask           : 255.255.0.0
Broadcast-Ip-Addr    : 13.13.255.255
Default-Router       : N/A
Primary-Dns          : 13.13.254.254
Secondary-Dns        : 13.13.254.253

ServerLeaseStart     : 12/24/2006 23:44:07
ServerLastRenew      : 12/24/2006 23:44:07
ServerLeaseEnd       : 12/24/2006 23:45:07
Session-Timeout      : 0d 00:01:00
DHCP Server Addr     : N/A

Persistent Relay Agent Information
  Circuit Id         : ancstb6_Dut-A|13|ies-13-13.13.1.1|0|13
  Remote Id          : stringtest
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-_Dut-A#
```

Routed CO Output Example

```
A:ALA-_Dut-A# show service id 13 dhcp lease-state
===============================================================================
DHCP lease state table, service 13
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
13.13.40.1      00:00:00:00:00:13 1/1/1:13            00h00m58s  Radius
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-_Dut-A#


A:ALA-_Dut-A# show service id 13 dhcp lease-state detail
===============================================================================
DHCP lease states for service 13
===============================================================================
Service ID           : 13
IP Address           : 13.13.40.1
Mac Address          : 00:00:00:00:00:13
Subscriber-interface : ies-13-13.13.1.1
Group-interface      : intf-13
SAP                  : 1/1/1:13
Remaining Lifetime   : 00h00m58s
Persistence Key      : N/A

Sub-Ident            : "Belgacom"
Sub-Profile-String   : "ADSL GO"
SLA-Profile-String   : "BE-Video"
Lease ANCP-String    : ""

Sub-Ident origin     : Radius
Strings origin       : Radius
Lease Info origin    : Radius

Ip-Netmask           : 255.255.0.0
Broadcast-Ip-Addr    : 13.13.255.255
Default-Router       : N/A
Primary-Dns          : 13.13.254.254
Secondary-Dns        : 13.13.254.253

ServerLeaseStart     : 12/24/2006 23:48:23
ServerLastRenew      : 12/24/2006 23:48:23
ServerLeaseEnd       : 12/24/2006 23:49:23
Session-Timeout      : 0d 00:01:00
DHCP Server Addr     : N/A

Persistent Relay Agent Information
  Circuit Id         : ancstb6_Dut-A|13|intf-13|0|13
  Remote Id          : stringtest
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-_Dut-A#
```

Wholesaler/Retainer Output Example

```
A:ALA-_Dut-A# show service id 2000 dhcp lease-state detail
===============================================================================
DHCP lease states for service 2000
===============================================================================
-------------------------------------------------------------------------------
```

```
Wholesaler 1000 Leases
-------------------------------------------------------------------------------
Service ID           : 1000
IP Address           : 13.13.1.254
Mac Address          : 00:00:00:00:00:13
Subscriber-interface : whole-sub
Group-interface      : intf-13
Retailer             : 2000
Retailer If          : retail-sub
SAP                  : 1/1/1:13
Remaining Lifetime   : 00h09m59s
Persistence Key      : N/A

Sub-Ident            : "Belgacom"
Sub-Profile-String   : "ADSL GO"
SLA-Profile-String   : "BE-Video"
Lease ANCP-String    : ""

Sub-Ident origin     : Retail DHCP
Strings origin       : Retail DHCP
Lease Info origin    : Retail DHCP

Ip-Netmask           : 255.255.0.0
Broadcast-Ip-Addr    : 13.13.255.255
Default-Router       : N/A
Primary-Dns          : N/A
Secondary-Dns        : N/A

ServerLeaseStart     : 12/25/2006 00:29:41
ServerLastRenew      : 12/25/2006 00:29:41
ServerLeaseEnd       : 12/25/2006 00:39:41
Session-Timeout      : 0d 00:10:00
DHCP Server Addr     : 10.232.237.2

Persistent Relay Agent Information
  Circuit Id         : 1/1/1:13
  Remote Id          : stringtest
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-_Dut-A#
```

statistics

Syntax **statistics** [**sap** *sap-id*
statistics [**sdp** *sdp-id*:*vc-id*]
statistics [**interface** *interface-name*]

Context show>service>id>dhcp

Description Displays DHCP statistics information.

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]

frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp-id — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Values 1 — 4294967295

interface *interface-name* — Displays information for the specified IP interface.

summary

Syntax **summary**

Context show>service>id>dhcp

Description Displays DHCP configuration summary information.

Output **Show DHCP Summary Output —** The following table describes the output fields for DHCP summary.

Label	Description
Interface Name	Name of the router interface.
Arp Populate	Specifies whether or not ARP populate is enabled.
Used/Provided	Used — The number of lease-states that are currently in use on a specific interface, that is, the number of clients on that interface got an IP address by DHCP. This value is always less than or equal to the 'Provided' field.
	Provided — The lease-populate value that is configured for a specific interface.
Info Option	Indicates whether Option 82 processing is enabled on the interface.
Admin State	Indicates the administrative state.

Sample Output

```
A:ALA-49# show service id 88 dhcp summary
===============================================================================
DHCP Summary, service 88
===============================================================================
Interface Name                   Arp      Used/                Info    Admin
  SapId/Sdp                      Populate Provided             Option  State
-------------------------------------------------------------------------------
Sector A                         No       0/0                  Keep    Up
  sap:7/1/1.2.2                           0/0
grp-if                           No       0/1                  Keep    Down
  sap:2/2/2:0                             0/1
test                             No       0/0                  Keep    Up
  sap:10/1/2:0                            0/0
-------------------------------------------------------------------------------
Interfaces: 3
===============================================================================
A:ALA-49#
```

gsmp

Syntax **gsmp**

Context show>service>id

Description This command displays 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                           Neigbor           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                           Neigbor           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                           Neigbor           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), neigbor 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), neigbor 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
===============================================================================
ANCP-String                                                  Assoc.
State
-------------------------------------------------------------------------------
7330-ISAM-E47 atm 1/1/01/01:19425.64048                      ANCP    Up
-------------------------------------------------------------------------------
Number of ANCP-Strings : 1
===============================================================================
A:active>show>service>id>gsmp#
```

host

Syntax **host**

Context show>service>id

Description Displays static hosts configured for this IES service.

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

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.

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* — Specifies the physical port identifier portion of the SAP definition.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	

qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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  5/2/3:0       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 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:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d

x: [0 — FFFF]H
d: [0 — 255]D

ip-int-name — 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 subscriber interfaces.

Values group, subscriber

detail — Displays detailed IP interface information.

Default IP interface summary output.

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

Label	Description
Interface-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 desired state of the interface.
Opr	The operating state of the interface.
Interface	
If Name	The name used to refer to the interface.
Admin State	The desired state of the interface.
Oper State	The operating state of the interface.
IP Addr/mask	Specifies the IP address/IP subnet/broadcast address of the interface.
Details	
If Index	The index corresponding to this IES interface. The primary index is 1; i.e., all IES 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.

Label	Description
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-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     7/1/1.2.2
   -                                                             -
test                             Up          Down/Down   IES     10/1/2:0
   1.1.1.1/30                                                    n/a
   1.1.1.1/30                                                    n/a
   1.1.2.1/30                                                    n/a
test27                           Up          Up/--       IES Sub subscriber
   192.168.10.21/24                                              n/a
grp-if                           Up          Down/--     IES Grp 2/2/2
-------------------------------------------------------------------------------
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.

Label	Description
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#
```

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.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128

bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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 operating state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP.
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.

Label	Description
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.
Forwarding Engine Stats	
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.
Off. Uncolor	The number of uncolored packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Queueing Stats (Ingress QoS Policy)	
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.
Queuing Stats (Egress QoS Policy)	
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.

Label	Description
Alarm Cell Han-dling	The OAM operational status of the VCL.
AAL-5 Encap	The AAL-5 encapsulation type.
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 7/1/1.2.2
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 88
SAP                : 7/1/1.2.2                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#
```

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.

subscriber-hosts

Syntax **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-address*[/*mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**detail**]

Context show>service>id

Description Displays subscriber host information.

Parameters **sap** *sap-id* — Displays the specified subscriber host SAP information.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200

qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

ip-address/mask — 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.

Values Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /30 subnets).
mask: 1 — 32

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

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

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

detail — Displays detailed information.

statistics

Syntax **statistics** [*ip-int-name* | *ip-address*]

Context show>router>dhcp

Description Display statistics for DHCP relay and DHCP snooping.

If no IP address or interface name is specified, then all configured interfaces are displayed.

If an IP address or interface name is specified, then only data regarding the specified interface is displayed.

Parameters *ip-int-name* | *ip-address* — Displays statistics for the specified IP interface.

Output **Show DHCP Statistics Output —** The following table describes the output fields for DHCP. statistics.

Label	Description
Received Packets	The number of packets received from the DHCP clients.

Label	Description
Transmitted Packets	The number of packets transmitted to the DHCP clients.
Received Malformed Packets	The number of malformed packets received from the DHCP clients.
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients.
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded.
Client Packets Relayed	The number of packets received from the DHCP clients that were forwarded.
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped.
Server Packets Discarded	The number of packets received from the DHCP server that were discarded.
Server Packets Relayed	The number of packets received from the DHCP server that were forwarded.
Server Packets Snooped	The number of packets received from the DHCP server that were snooped.

Sample Output

```
*A:ALA-1# show router dhcp statistics
===============================================================
DHCP Global Statistics
===============================================================
Rx Packets                        : 0
Tx Packets                        : 0
Rx Malformed Packets              : 0
Rx Untrusted Packets              : 0
Client Packets Discarded          : 0
Client Packets Relayed            : 0
Client Packets Snooped            : 0
Server Packets Discarded          : 0
Server Packets Relayed            : 0
Server Packets Snooped            : 0
===============================================================
*A:ALA-1#
```

summary

Syntax **summary**

Context show>router>dhcp

Description Display the status of the DHCP relay and DHCP snooping functions on each interface.

Output **Show DHCP Summary Output —** The following table describes the output fields for DHCP summary.

Label	Description
Interface Name	Name of the router interface.
SapId/Sdp	Specifies the associated SAP ID or SDP ID.
Arp Populate	Specifies whether or not ARP populate is enabled.
Used/Provided	Used – The number of lease-states that are currently in use on a specific interface, that is, the number of clients on that interface got an IP address by DHCP. This value is always less than or equal to the 'Provided' field.
	Provided – The lease-populate value that is configured for a specific interface.
Info Option	Indicates whether Option 82 processing is enabled on the interface.
Admin State	Indicates the administrative state.

Sample Output

```
A:ALA-49# show router dhcp summary
===============================================================================
DHCP Summary (Router: Base)
===============================================================================
Interface Name                  Arp      Used/                Info    Admin
  SapId/Sdp                     Populate Provided             Option  State
-------------------------------------------------------------------------------
Sector A                        No       0/0                  Keep    Up
  sap:7/1/1.2.2                          0/0
grp-if                          No       0/1                  Keep    Down
ies-test                        No       0/0                  Keep    Up
  sap:9/1/2:0/500                        0/0
test                            No       0/0                  Keep    Up
  sap:10/1/2:0                           0/0
test1                           No       0/0                  Keep    Up
  sap:7/1/1.1.2                          0/0
test2                           No       0/0                  Keep    Up
  sap:7/1/1.2.1                          0/0
testA                           No       0/0                  Keep    Up
  sap:7/1/3.1.1                          0/0
testB                           No       0/0                  Keep    Up
  sap:7/1/5.1.1                          0/0
to-HQ                           No       0/0                  Keep    Up
  sdp:spoke-2:1001                       0/0
to-web                          No       0/0                  Keep    Up
  sap:2/1/10:50                          0/0
-------------------------------------------------------------------------------
Interfaces: 10
===============================================================================
A:ALA-49#
```

Clear Commands

dhcp

Syntax **dhcp**

Context clear>router>dhcp

Description This command enables the context to clear DHCP parameters.

dhcp6

Syntax **dhcp6**

Context clear>router>dhcp6

Description This command enables the context to clear DHCP6 parameters.

statistics

Syntax **statistics** [*ip-int-name* | *ip-address*]

Context clear>router>dhcp

Description Clears DHCP statistics.

id

Syntax **id** *service-id*

Context clear>service
clear>service>statistics

Description This command clears parameters for a specific service.

Parameters *service-id* — The ID that uniquely identifies the service to clear.

interface

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.

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 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.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type*-*slot/mda*.*bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type*-*bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200

qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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.

Values 1 — 4294967295

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.

wholesalers

Syntax **wholesalers**

Context show>service>id

Description This command displays service wholesaler information.

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

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.

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 /30 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 *sap-id* — Clears the specified lease state SAP information.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port*.*channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type*-*slot/mda*.*bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type*-*bpgrp-num*	

	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp-id — The specified SDP to be cleared.

Values 1 — 17407

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

Values 1 — 4294967295

lease-state

Syntax **lease-state** [**ip-address** *ipv6-address*/*prefix-length*] [**mac** *ieee-address*]

Context clear>service>id>dhcp6

Description Clears DHCP6 lease state information for this service.

Parameters **ip-address** *ipv6-address/prefix-length* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /30 subnets).

Values

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x [0 — FFFF]H d [0 — 255]D
prefix-length	1 — 128

mac *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.

statistics

Syntax **statistics** [*ip-int-name* | *ipv6-address*]

Context clear>router>dhcp6

Description Clears DHCP6 statistics.

Parameters *ip-int-name* — Specifies the IP interface name up to 32 characters inlength.

ip-address *ipv6-address/prefix-length* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /30 subnets).

Values		
ipv6-address	x:x:x:x:x:x:x:x	(eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d	
	x [0 — FFFF]H	
	d [0 — 255]D	
prefix-length	1 — 128	

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id* | **interface** *ip-address* | *ip-int-name*]

Context clear>service>id>dhcp

Description Clears DHCP statistics.

Parameters **sap** *sap-id* — Clears the specified SAP information.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8

	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp *sdp-id* — The specified SDP to be cleared.

Values 1 — 17407

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

Values 1 — 4294967295

interface *ip-address* — The interface IP address.

ip-int-name — The interface name.

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

RFC2547bis is an extension to the original RFC 2547, which details a method of distributing routing information and forwarding data to provide a Layer 3 Virtual Private Network (VPN) service to end customers.

Each Virtual Private Routed Network (VPRN) consists of a set of customer sites connected to one or more PE routers. Each associated PE router maintains a separate IP forwarding table for each VPRN. Additionally, the PE routers exchange the routing information configured or learned from all customer sites via MP-BGP peering. Each route exchanged via the MP-BGP protocol includes a Route Distinguisher (RD), which identifies the VPRN association.

The service provider uses BGP to exchange the routes of a particular VPN among the PE routers that are attached to that VPN. This is done in a way which ensures that routes from different VPNs remain distinct and separate, even if two VPNs have an overlapping address space. The PE routers distribute routes from other CE routers in that VPN to the CE routers in a particular VPN. Since the CE routers do not peer with each other there is no overlay visible to the VPN's routing algorithm.

When BGP distributes a VPN route, it also distributes an MPLS label for that route. On a SR-Series, a single label is assigned to all routes in a VPN.

Before a customer data packet travels across the service provider's backbone, it is encapsulated with the MPLS label that corresponds, in the customer's VPN, to the route which best matches the packet's destination address. The MPLS packet is further encapsulated with either another MPLS label 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 52 displays a VPRN network diagram.

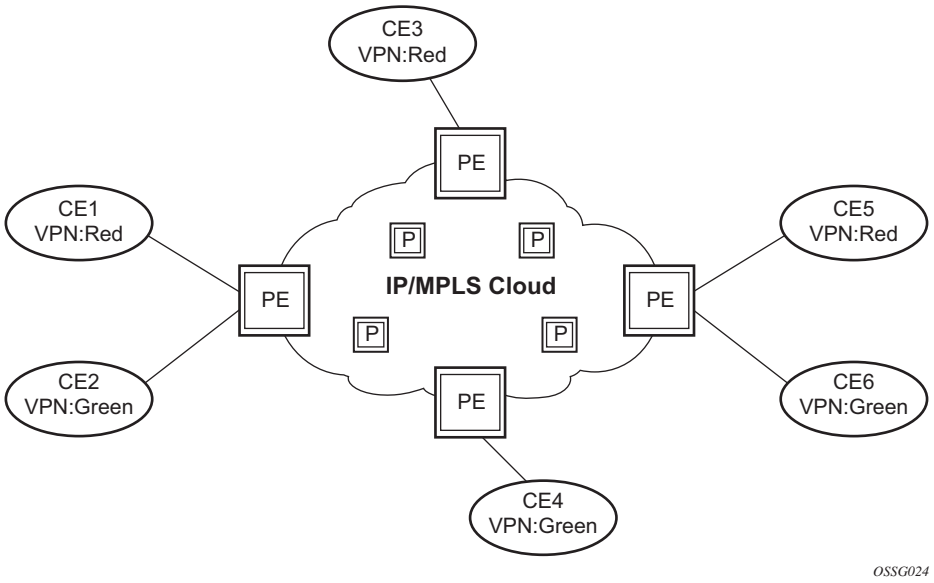

Figure 1: Virtual Private Routed Network

Routing Prerequisites

RFC2547bis requires the following features:

- Multi-protocol extensions
- LDP support
- Extended BGP community support
- BGP capability negotiation
- Parameters defined in RFC 2918, *BGP Route Refresh* and RFC 2796, *Route Reflector*
- 4-byte autonomous system (AS) number

Tunneling protocol requirements are as follows:

- RFC2547, *BGP/MPLS VPNs*, recommends implementing Label Distribution Protocol (LDP) to set up a full mesh of LSPs based on the IGP
- MPLS RSVP-TE tunnels can be used instead of LDP
- Generic Router Encapsulation (GRE) tunnels can also be alternatively used

BGP Support

BGP is used with BGP extensions mentioned in Routing Prerequisites on page 857 to distribute VPRN routing information across the service provider's network.

BGP was initially designed to distribute IPv4 routing information. Therefore, multi-protocol extensions and the use of a VPN-IPv4 address were created to extend BGP's ability to carry overlapping routing information. A VPN-IPv4 address is a 12-byte value consisting of the 8-byte route distinguisher (RD) and the 4-byte IPv4 IP address prefix. The RD must be unique within the scope of the VPRN. This allows the IP address prefixes within different VRFs to overlap.

Route Distinguishers

The route distinguisher (RD) is an 8-byte value consisting of 2 major fields, the Type field and value field. The type field determines how the value field should be interpreted. The TiMOS implementation must support the three (3) type values as defined in the internet draft.

Figure 2: 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.

Route Reflector

Per RFC2547bis the use of Route Reflectors must be supported in the service provider core. The support must use multiple sets of route reflectors for different types of BGP data, including IPv4 and VPN-IPv4 (In the future multicast and IPv6).

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

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.

Routes belonging to a VPRN, must use the protocol owner, VPN-IPv4 to denote that it is an VPRN route. This can be used within the route policy match criteria.

VPRN Features

This section describes various of the general 7750 SR service features and any special capabilities or considerations as they relate to VPRN services.

Subscriber Interfaces

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

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

For proper operation, each subscriber subnet associated with the SRRP instance must have a gw-address defined. The SRRP instance cannot be activated (no shutdown) unless each subscriber subnet associated with the group IP interface has an SRRP gateway IP address. Once the SRRP instance is activated, new subscriber subnets cannot be added without a corresponding SRRP gateway IP address. Table 2.2 describers how the SRRP instance state is used to manage access to subscriber hosts associated with the group IP interface.

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

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

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

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

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

Table 30 lists the SRRP's state effect on subscriber hosts associated with group IP interfaces.

Table 1: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface

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

Table 1: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface (Continued)

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

Table 1: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface (Continued)

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

Table 1: SRRP State Effect on Subscriber Hosts Associated with Group IP Interface (Continued)

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

SRRP Messaging

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

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

The vr-id field has been expanded to support an SRRP instance ID of 32 bits.

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

Unlike VRRP, only two nodes may participate in an SRRP instance due the explicit association between the SRRP instance group IP interface, the associated redundant IP interface and the multi-chassis synchronization (MCS) peering. Since only two nodes are participating, the VRRP skew timer is not utilized when waiting to enter the master state. Also, SRRP always preempts

when the local priority is better than the current master and the backup SRRP instance always inherits the master's advertisement interval from the SRRP advertisement messaging.

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

The SRRP advertisement message is always evaluated to see if it has higher priority than the SRRP advertisement that would be sent by the local node. If the advertised priority is equal to the current local priority, the source IP address of the received SRRP advertisement is used as a tie breaker. The node with the lowest IP address is considered to have the highest priority.

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

SRRP and Multi-Chassis Synchronization

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

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

- The SRRP instance MCS key
- Containing service type and ID
- Containing subscriber IP interface name
- Subscriber subnet information

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

SRRP Instance

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

- SRRP Instance MCS Key on page 867
- Containing Service Type and ID on page 867
- Containing Subscriber IP Interface Name on page 868
- Subscriber Subnet Information on page 868

SRRP Instance MCS Key

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

Containing Service Type and ID

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

Containing Subscriber IP Interface Name

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

Subscriber Subnet Information

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

Containing Group IP Interface Information

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

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

Remote Redundant IP Interface Mismatch

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

Remote Sending Redundant IP Interface Unavailable

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

Remote SRRP Advertisement SAP Non-existent

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

Remote Sending Local Receive SRRP Advertisement SAP Unavailable

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

Local and Remote Dual Master Detected

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

Subscriber Subnet Owned IP Address Connectivity

In order for the network to reliably reach the owned IP addresses on a subscriber subnet, the owning node must advertise the IP addresses as /32 host routes into the core. This is important since the subscriber subnet is advertised into the core by multiple routers and the network will follow the shortest path to the closest available router which may not own the IP address if the /32 is not advertised within the IGP.

Subscriber Subnet SRRP Gateway IP Address Connectivity

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

Receive SRRP Advertisement SAP and Anti-Spoof

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

SAPs

Encapsulations

The following SAP encapsulations are supported on the 7750 SR VPRN service:

- Ethernet null
- Ethernet dot1q
- SONET/SDH IPCP
- SONET/SDH ATM
- ATM - LLC SNAP or VC-MUX
- Cisco HDLC
- Frame Relay

ATM SAP Encapsulations for VPRN Services

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

The following ATM encapsulation and transport modes are supported:

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

QoS Policies

When applied to 7750 SR VPRN services, service ingress QoS policies only create the unicast queues defined in the policy. The multipoint queues are not created on the service.

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 L2 or L3 criteria can be used in the QoS policies for traffic classification in an VPRN.

ATM traffic descriptor profiles define the traffic contract in the forward direction (ingress) and the traffic contract in the backward direction (egress).

Filter Policies

Only IP filter policies can be applied to VPRN services.

CE to PE Routing Protocols

The 7750 SR VPRN supports the following PE to CE routing protocols:

- BGP
- Static
- RIP
- OSPF

PE to PE Tunneling Mechanisms

The 7750 SR supports multiple mechanisms to provide transport tunnels for the forwarding of traffic between PE routers within the 2547bis network.

The 7750 SR 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 7750 SR 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.

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:

- Transportation of OSPF learned routes as OSPF externals. This feature uses OSPF as the protocol between the PE and CE routers, however instead of transporting the OSPF LSA information across the IP-VPN, the OSPF routes are "imported" into MP-BGP as AS externals. As a result, other OSPF attached VPRN sites on remote PEs receive these via type 5 LSA.

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 Service Distribution Points (SDPs) on page 41 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 54 depicts traffic terminating on a specific IES or VPRN service that is identified by the *sdp-id* and VC label present in the the service packet.

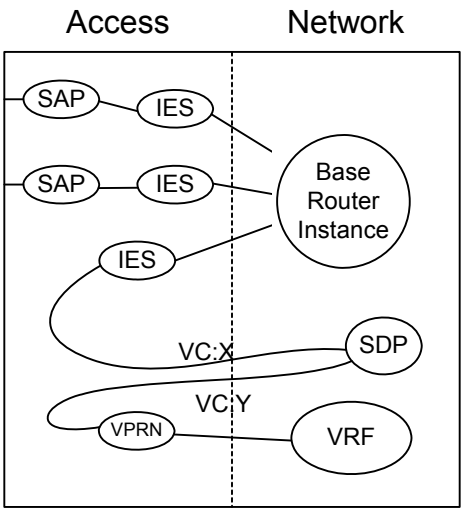

Figure 3: SDP-ID and VC Label Service Identifiers

Multicast in IP-VPN Applications

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 in the future, video delivery systems.

Implementation of the draft-rosen-vpn-mcast, *Multicast in MPLS/BGP 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.

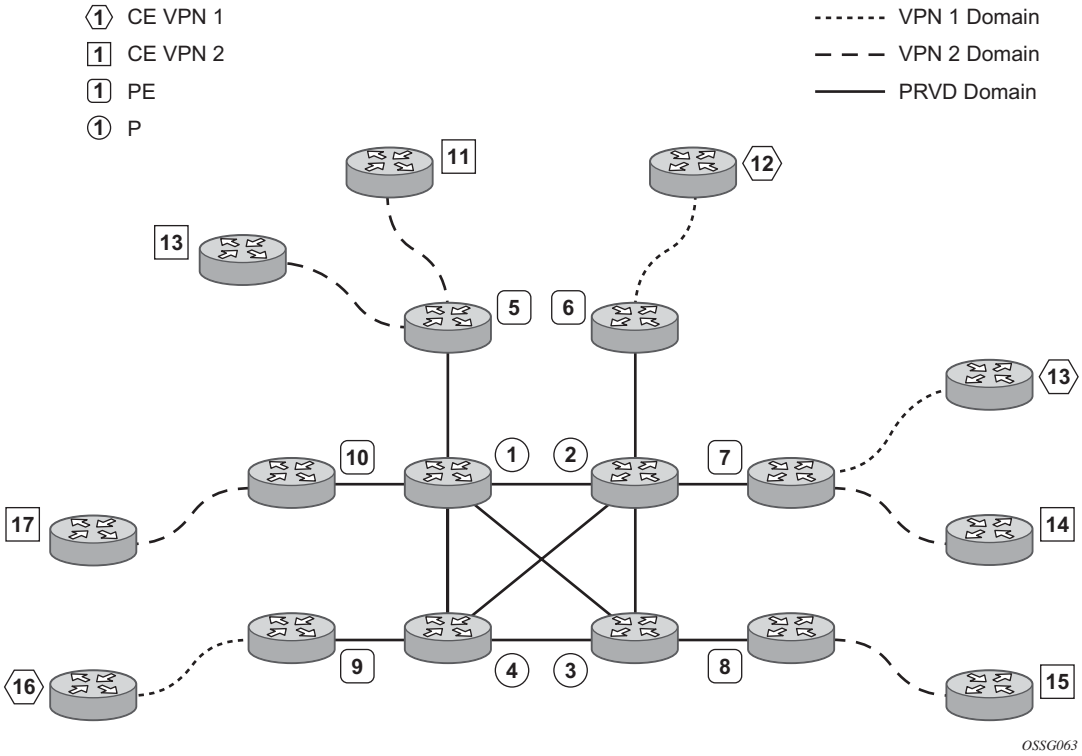

Figure 4: Multicast in IP-VPN Applications

Figure 55 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

PIM-SM for default-mdts and PIM-SSM for data-mdts in the core network are supported by Draft Rosen. In the customer network, both PIM-SM and PIM-SSM are supported.

Cflowd for IP-VPNs (VPRNs)

The cflowd feature allows service providers to collection IP flow data within the context of a VPRN. This data can 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 should not be used for billing purposes. Existing queue counters are designed for this purpose and provide very accurate per bit accounting records.

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:

List of Commands

Table 31 lists all the service configuration commands indicating the configuration level at which each command is implemented with a short command description. VPRN services are configured in the `config>service>vprn` context. The command list is organized in the following task-oriented manner:

- Configure a VPRN service
- Configure VPRN service parameters
  - → Configure VPRN IGMP parameters
  - → Configure VPRN spoke SDP parameters
- Configure VPRN subscriber interface parameters
  - → Configure VPRN subscriber interface group interface parameters
- Configure VPRN service PIM parameters
  - → Configure VPRN service PIM interface parameters
- Configure VPRN interface parameters
  - → Configure VPRN interface SAP parameters
  - → Configure SAP subscriber management parameters
  - → Configure VPRN interface SAP ATM parameters
  - → Configure VPRN interface DHCP parameters
  - → Configure VPRN interface VRRP parameters
- Configure VPRN BGP parameters
  - → Configure group level VPRN BGP group parameters
  - → Configure group level VPRN BGP neighbor parameters
- Configure VPRN RIP parameters
  - → Configure VPRN RIP group parameters
  - → Configure VPRN RIP neighbor parameters
- Configure VPRN OSPF parameters

Table 1: CLI Commands to Configure VPRN Service Parameters

Command	Description	Page
Configure a VPRN service		903
`config>service>vprn`		
`service-id`	Specifies a unique service identification number identifying the service in the service domain.	954

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
customer	Specifies the existing customer ID number to be associated with the service.	954
Configure VPRN service parameters		
config>service>vprn		
description	Specifies a text string describing the service.	952
auto-bind	Specifies the automatic binding type for the SDP assigned to this service.	955
autonomous-system	Defines the autonomous system (AS) to be used by this VPN routing/forwarding (VRF).	955
ecmp	Enables ECMP and configures the number of routes for path sharing.	955
gsmp	Enables the context to configure GSMP connections maintained in this service.	956
group	Specifies a GSMP name.	956
ancp	Configures ANCP parameters for this GSMP group.	956
dynamic-topology-discover	Enables the ANCP dynamic topology discovery capability.	956
oam	Specifies GSMP ANCP OAM negotiation at the start of a GSMP connection.	957
hold-multiplier	Configures the hold-multiplier for the GSMP connection in this group.	957
keep-alive	Configures keepalive values for the GSMP connection in this group.	957
neighbor	Adds or removes a neighbor from this GSMP group.	957
local-address	Configures the source IP address.	958
priority-marking	Configures the type of poriority marking.	958
igmp	Enables the context to configure IGMP connections maintained in this service.	958
maximum-routes	Specifies the maximum number of routes that can be held within a VPN routing/forwarding (VRF) context.	965
mc-maximum-routes	Specifies the maximum number of multicast routes that can be held within a VPN routing/forwarding (VRF) context.	966
route-distinguisher	Sets the identifier attached to routes the VPN belongs to.	967
router-id	Sets the router ID for a specific VPRN context.	967
snmp-community	Sets the SNMP community name to be used with the associated VPRN instance.	967
static-route	Creates static route entries for network and access routes.	968
type	Designates the type of VPRN instance being configured for hub and spoke topologies.	970

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
vrf-export	Specifies the export policies to control routes exported from the local VPN routing/forwarding (VRF) to remote PE routers (for example, through MP-BGP).	970
vrf-import	Specifies the import policies to control routes imported to the local VPN routing/forwarding (VRF) from remote PE Routers (via MP-BGP).	971
vrf-target	Facilitates a simplified method to configure the route target to be added to advertised routes to remote PEs or compare against received routes from remote PEs.	971
no shutdown	Administratively enables the VPRN service.	951
Configure redundant interface parameters		
config>service>vprn		
redundant-interface	Configures a subscriber interface.	972
address	Assigns an IP address/IP subnet format and a remote IP to the interface.	972
description	A text string that describes the redudant interface.	952
no shutdown	Administratively enables the interface.	951
Configure VPRN IGMP parameters		
config>service>vprn>igmp		
interface	Configures the IGMP interface.	958
import	Imports a policy to filter IGMP packets.	959
max-groups	Configures the maximum number of groups for this interface.	959
mcac	Configures multicast CAC policy and constraints for this interface.	959
policy	Configures mulitcast CAC policy name.	
mc-constraints	Configures multicast CAC constraints.	960
level	Configures levels and their associated bandwidth for multicast CAC policy on this interface.	960
number-down	Configures the number of ports down along with level for multicast CAC policy on this interface.	960
shutdown	Administratively enable/disable constraint for multicast cac policy.	951
policy	Configures the mulitcast CAC policy name.	960
unconstrained-bw	Configures unconstrained-bw for multicast CAC policy on this interface.	961
static	Adds IGMP static group membership.	961
group	Adds a static multicast group either as a (*,G) or one or more (S,G) records.	962

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
source	Adds a static multicast source.	962
starg	Adds a static starg entry.	962
subnet-check	Enables or disables local subnet checking for IGMP.	963
version	Configures the version of IGMP.	963
query-interval	Configures the frequency at which Host-Query packets are transmitted.	963
query-last-member-interval	Configures the frequency at which Group-Specific-Query packets are transmitted.	964
query-response-interval	Configures the time to wait to receive a response to the Host-Query message from the host.	964
robust-count	Configures the robust count.	964
ssm-translate	Enables the context to add or remove ssm-translate group ranges.	965
grp-range	Adds an SSM translate group range entry.	965
source	Adds the source address for the SSM translate group range.	965
Configure VPRN spoke SDP parameters		
config>service>vprn		
spoke-sdp	Binds a service to an existing Service Distribution Point (SDP).	973
no shutdown	Administratively enables the spoke SDP binding.	951
Configure VPRN subscriber interface parameters		
config>service>vprn>sub-if		
address	Assigns an IP address, IP subnet, and broadcast address format to an VPRN IP router interface.	995
authentication-policy	Assigns a RADIUS authentication policy to the interface.	1010
dhcp	Configures DHCP parameters for the subscriber interface	987
group-interface	Enables the context to configure a group interface.	996
no shutdown	Administratively enables the subscriber interface.	951
Configure VPRN subscriber interface group interface parameters		
config>service>vprn>sub-if>grp-if		
arp-populate	Disables dynamic learning of ARP entries.	1009
arp-timeout	Configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table.	1010
authentication-policy	Assigns a RADIUS authentication policy to the interface.	1010
description	Creates a text description stored in the configuration file for a configuration context.	952
dhcp	Configures DHCP parameters for this interface	987
host-connectivity	Enables host connectivity verification for all hosts on this interface.	994

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
local-proxy-arp	Enables local proxy ARP on the interface.	981
mac	Assigns a specific MAC address to a VPRN interface.	982
proxy-arp-policy	Configures a proxy ARP policy for the interface.	982
remote-proxy-arp	Enables remote proxy ARP on the interface.	983
tos-marking-state	Changes the default trusted state to a non-trusted state.	985
Configure IES group interface SRRP parameters		
config>service>ies>sub-if>grp-if>srrp		
gw-mac	Overrides the default SRRP gateway MAC address used by the SRRP instance.	1035
keep-alive-interval	defines the interval between SRRP advertisement messages sent when operating in the master state.	1036
message-path	defines a specific SAP for SRRP in-band messaging.	1036
policy	associates one or more VRRP policies with the SRRP instance.	1038
priority	overrides the default base priority for the SRRP instance.	1038
no shutdown	Administratively enables the entity.	951
Configure VPRN service PIM parameters		906
config>service>vprn>pim		
apply-to	Creates a PIM interface with default parameters.	1048
import	Specifies the import route policy to be used for determining which routes are accepted from peers.	1048
mdt	Enables the context for a multicast distribution tree (MDT) to carry multicast traffic from customer sites associated with the multicast domain.	1053
data	Configures a pool of addresses that can be used to generate data only MDT tunnels.	1054
data-delay-interval	Configures the data delay interval in seconds.	1054
data-threshold	Configures threshold for a group prefix.	1054
default	Configures a default multicast distribution tree (MDT) group address used by the core instance of PIM to identify multicast traffic for this VPRN instance.	1055
join-tlv-packing-disable	Enables the packing of multiple MDT join TLVs.	1055
non-dr-attract-traffic	Specifies whether the router should ignore the designated router state and attract traffic even when it is not the designater router.	1055
hello-interval	Configures the frequency at which PIM Hello messages are transmitted on this interface.	1050

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
hello-multiplier	Configures the multiplier to determine the holdtime for a PIM neighbor.	1050
improved-assert	Enables improved assert processing on this interface.	1051
max-groups	Configures the maximum number of groups for which PIM can have downstream state based on received PIM Joins on this interface.	1051
tracking-support	Sets the the T bit in the LAN Prune Delay option of the Hello Message.	1053
ssm-groups	Enables access to the context to enable a source-specific multicast (SSM) configuration instance.	1062
group-range	Configures the group address or range of group addresses for which this router can be the rendezvous point (RP).	1059
no shutdown	Administratively enables the VPRN PIM instance.	951
Configure VPRN service PIM interface parameters		916
config>service>vprn>pim>interface		
bsm-check-rtr-alert	Enables the checking of router alert option in the bootstrap messages received on this interface.	1050
hello-interval	Configures the frequency at which PIM Hello messages are transmitted on this interface.	1050
multicast-senders	Configures the way subnet matching is done for incoming data packets on this interface.	1051
priority	Sets the priority value to become the rendezvous point (RP) that is included in bootstrap messages sent by the router.	1052
tracking-support	Sets the the T bit in the LAN Prune Delay option of the Hello Message.	1053
no shutdown	Administratively enables the VPRN PIM interface.	951
config>service>vprn>pim>rp		
anycast	Configures a PIM anycast protocol instance for the RP being configured.	1056
rp-set-peer	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.	1057
bootstrap-export	Exports policies to control the flow of bootstrap messages from the RP. Up to five policies can be defined.	1057
bootstrap-import	Imports policies to control the flow of bootstrap messages into the RP.	1057
bsr-candidate	Enables the context to configure a local rendezvous point (RP) of a PIM protocol instance.	1058
rp-candidate	Enables the context to configure the candidate rendezvous point (RP) parameters.	1059

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
static	Enables access to the context to configure a static rendezvous point (RP) of a PIM-SM protocol instance.	1060
address	Configures the static rendezvous point (RP) address.	1060
hash-mask-len	Configure the length of a mask that is to be combined with the group address before the hash function is called.	1058
priority	Configures the priority used to become the rendezvous point (RP).	1058
group-range	Configures the group address or range of group addresses for which this router can be the rendezvous point (RP).	1059
holdtime	Configures the length of time neighboring routers consider this router to be up.	1059
group-prefix	Configures the multicast group-address prefix which contains multicast group-addresses that are imbedded in the join or prune packet as a filter criterion.	1061
no shutdown	Administratively enables the RP BSR or rendezvous point (RP) candidate.	951
Configure VPRN interface parameters		915
config>service>vprn>interface		
address	Assigns the primary IP address, IP subnet, and broadcast address format to an IP interface.	978
allow-directed-broadcast	Enables directed broadcast forwarding out of the IP interface.	980
arp-populate	Enables populating static and dynamic hosts into the system ARP cache.	1009
arp-timeout	Configures the minimum time, in seconds, an ARP entry learned on the IP interface is stored in the ARP table.	1010
authentication-policy	Configures the authentication policy for this policy.	1011
bfd	Specifies the bi-directional forwarding parameters for the IP interface.	980
cflowd	Configures cflowd collection and analysis on the (routeable) interface.	981
description	Creates a text description for the interface.	952
dhcp	Enables the context to configure DHCP parameters.	987
host	Configures one or more hosts on this SAP.	1013
local-proxy-arp	Enables local proxy ARP on an interface.	981
loopback	Configures the interface as a loopback interface. This command cannot be issued if a SAP is defined on the interface.	981
mac	Assigns a specific MAC address to an IP interface.	982
proxy-arp	Configures a proxy ARP policy for the interface.	982

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
secondary	Assigns a secondary IP address, IP subnet, and broadcast address format to an IP interface.	983
static-arp	Configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the router instance.	984
no shutdown	Enables the VPRN interface.	951
Configure VPRN interface SAP parameters		917
config>service>vprn>interface>sap		
accounting-policy	Specifies the accounting policy to apply to the SAP.	1004
anti-spoof	Enables anti-spoof filtering and optionally changes the anti-spoof matching type for the interface.	1009
collect-stats	Enables the collection of accounting and statistical data for the SAP or network port	1004
description	Creates a text description stored in the configuration file for the SAP.	952
ingress	Enables a context to configure ingress SAP Quality of Service (QoS) policies and filter policies.	1006
multi-service-site	Creates a new customer site or edits an existing customer site with the *customer-site-name* parameter.	1022
tod-suite	Applies a time-based policy (filter or QoS policy) to the SAP.	1003
Configure SAP egress parameters		
config>service>vprn>if>sap>egress		
filter	Associates an IP filter policy or MAC filter policy with an egress SAP or IP interface.	1016
qos	Associates a Quality of Service (QoS) policy with an egress SAP or IP interface.	1020
queue-override	Enables the context to configure override values for the specified SAP egress QoS queue.	1023
queue	Specifies the ID of the queue whose parameters are to be overridden.	1023
adaptation-rule	Specifies attributes of the specified queue's adaptation rule parameters.	1023
cbs	Specifies attributes of the specified queue's CBS parameters.	1026
high-prio-only	Specifies attributes of the specified queue's high-prio-only parameters.	1026
mbs	Specifies attributes of the specified queue's MBS parameters.	1027
rate	Specifies attributes of the queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.	1028
scheduler-override	Specifies the set of attributes whose values have been overridden via management on this virtual scheduler.	1029

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
scheduler	Defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler.	1029
scheduler-policy	Associates an existing scheduler policy to an egress scheduler used by SAP queues associated with this multi-service customer site.	1021
Configure SAP ingress parameters		
config>service>vprn>if>sap>ingress		
filter	Associates an IP or MAC filter policy with an ingress SAP or IP interface.	1016
match-qinq-dot1p	Configures filtering based on the p-bits in the top or bottom tag of a Q-in-Q encapsulated Ethernet frame.	1016
qos	Associates a Quality of Service (QoS) policy with an ingress SAP or IP interface.	1020
queue-override	Enables the context to configure override values for the specified SAP egress QoS queue.	1023
queue	Specifies the ID of the queue whose parameters are to be overridden.	1023
adaptation-rule	Overrides specific attributes of the specified queue's adaptation rule parameters.	1023
cbs	Overrides specific attributes of the specified queue's CBS parameters.	1026
high-prio-only	Overrides specific attributes of the specified queue's high-prio-only parameters.	1026
mbs	Overrides specific attributes of the specified queue's MBS parameters.	1027
rate	Overrides specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.	1028
scheduler-override	Specifies the set of attributes whose values have been overridden via management on this virtual scheduler.	1029
scheduler	Defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler.	1029
scheduler-policy	Associates an existing scheduler policy to an ingress scheduler used by SAP queues associated with this multi-service customer site.	1021
Configure VPRN interface SAP ATM parameters		917
config>service>vprn>interface>sap		
atm	Enables access to the context to configure ATM-related attributes.	1005
egress	Configures egress ATM attributes for the SAP.	1005
traffic-desc	Assigns an ATM traffic descriptor profile to a given context (for example a SAP).	1006
encapsulation	Configures RFC 2684 encapsulation for an ATM PVCC delimited SAP.	1005

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
ingress	Configures ingress ATM attributes for the SAP.	1006
oam	Enables the context to configure OAM functionality for a PVCC delimiting a SAP.	1006
alarm-cells	Configures AIS/RDI fault management on a PVCC.	1007
periodic-loopback	Enables periodic OAM loopbacks on this SAP.	1007
Configure VPRN interface DHCP parameters		
config>service>vprn>if>dhcp		
dhcp	Configures DHCP parameters for this interface	987
description	Specifies a text string description for DHCP on this interface.	952
gi-address	Configures the gateway interface address for the DHCP relay.	988
lease-populate	Enables dynamic host lease state management for VPRN IP interfaces.	988
option	Configures DHCP Option 82 processing for this interface.	1007
server	Configures the DHCP server IP address.	993
no shutdown	Enables DHCP on this interface.	951
trusted	Enables relaying of untrusted packets.	993
config>service>vprn>if>dhcp>option		
action	Configures the DHCP relay reforwarding policy action.	987
circuit-id	Enables sending of the interface index in the circuit-id suboption of the DHCP relay packet.	987
remote-id	Enables sending of the remote MAC address in the remote-id suboption of the DHCP relay packet.	990
To configure interface ICMP parameters:		904
config>service>vprn *service-id*>interface		
icmp	Configures ICMP parameters on an VPRN IP interface.	997
mask-reply	Enables responses to ICMP mask requests on the router interface.	997
redirects	Enables and configures the rate for ICMP redirect messages issued on the router interface.	997
unreachables	Enables and configures the rate for ICMP host and network destination unreachable messages issued on the router interface.	998
ttl-expired	Configures the rate ICMP TTL expired messages are issued by the IP interface.	998
Configure VPRN interface VRRP parameters		
config>service>vprn>interface>vrrp		

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
authentication-key	Sets/clears the simple text authentication key used for generating master VRRP advertisement messages and validating received VRRP advertisements.	1040
authentication-type	Configures the VRRP authentication: • VRRP Type 0 authentication provides no authentication. All compliant VRRP advertisement messages are accepted. • VRRP Type 1 authentication provides a simple password check on incoming VRRP advertisement messages. • VRRP Type 2 authentication provides an MD5 IP header authentication check on incoming VRRP advertisement messages.	1041
backup	Assigns virtual router IP addresses associated with the parental IP interface IP addresses. Owner instances do not create a routable IP interface address; it defines the existing parental IP interface IP addresses that will be advertised by the virtual router instance.	1043
mac	Sets an explicit MAC address to be used by the virtual router instance overriding the VRRP default derived from the VRID.	1043
master-int-inherit	Allows the master instance to dictate the master down timer (non-owner context only).	1043
message-interval	Configures the administrative advertisement message timer used by the master virtual router instance to send VRRP advertisement messages and to derive the master down timer as backup.	1044
ping-reply	Enables the non-owner master to reply to ICMP Echo Requests directed at the virtual router instance's IP addresses.	1044
preempt	Provides the ability of overriding an existing non-owner master to the virtual router instance.	1045
priority	Provides the ability to configure a specific priority value to the virtual router instance.	1045
no shutdown	Enables the VPRN interface VRRP instance.	951
ssh-reply	Enables the non-owner master to reply to SSH Requests directed at the virtual router instance's IP addresses.	1046
telnet-reply	Enables the non-owner master to reply to TCP port 23 Telnet Requests directed at the virtual router instance's IP addresses.	1047
traceroute-reply	Enables a non-owner master to reply to traceroute requests directed to the virtual router instance's IP addresses.	1047

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
Configure VPRN BGP parameters		909
config>service>vprn# bgp		
advertise-inactive	Enables the advertising of inactive BGP routers to other BGP peers.	1063
aggregator-id-zero	Enables setting the router ID to zero in the aggregator path attribute when BGP is aggregating routes.	1063
always-compare-med	Determines how the Multi-Exit Discriminator (MED) path attribute is used in the BGP route selection process.	1064
as-path-ignore	Determines whether the AS path is used in determining the best BGP route.	1064
authentication-key	Enables MD5 authentication and configures the authentication key.	1065
cluster	Configures the cluster ID for a route reflector server.	1066
connect-retry	Configures the BGP Connect Retry timer value in seconds.	1066
damping	Enables BGP route damping for learned routes as defined within the route policy.	1067
description	Creates a text description for the BGP instance.	952
disable-client-reflect	Disables the reflection of routes by the route reflector to the clients in the given peer group.	1067
disable-communities	Specifies BGP to disable sending communities.	1068
disable-fast-external-failover	Configures BGP fast external failover.	1068
export	Specifies the export route policy to be used for determining which routes are advertised to peers.	1068
hold-time	Configures the BGP Hold Time in seconds.	1070
ibgp-multipath	Allows BGP to use multiple BGP learned routes with different BGP nexthops and also, to use ECMP routes to resolve the BGP nexthop providing either load-balancing across the ECMP nexthops or install the BGP route with all of the ECMP nexthops.	1070
import	Specifies the import route policy to be used for determining which routes are accepted from peers.	1070
keepalive	Configures the BGP keepalive timer in seconds.	1071
local-as	Adds a BGP virtual AS number.	1072
local-preference	Enables setting the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute.	1073
loop-detect	Configures how the BGP peer session will handle loop detection in the AS Path.	1073

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
med-out	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.	1074
min-as-origination	Configures the minimum interval in seconds at which a given path attribute, originated by the local router, can be advertised to a peer.	1075
min-route-advertisement	Configures the minimum interval in seconds at which a given prefix can be advertised to a peer.	1075
multihop	Configures the TTL value entered in the IP header of packets sent to an EBGP peer multiple hops away.	1076
multipath	Enables BGP multipath. When multipath is enabled BGP will load share traffic across multiple links.	1076
preference	Configures the route preference for routes learned from the configured peer(s).	1078
remove-private	Enables removing the private AS numbers from the AS Path before advertising them to BGP peers.	1079
router-id	Configures the router ID to be used with this BGP instance.	967
no shutdown	Administratively enables BGP.	951
Configure group level VPRN BGP group parameters		910
config>service>vprn>bgp>group		
advertise-inactive	Enables the advertising of inactive BGP routers to other BGP peers.	1063
aggregator-id-zero	Enables setting the router ID to zero in the aggregator path attribute when BGP is aggregating routes.	1063
as-override	Replaces all instances of the peer's AS number with the local AS number in a BGP route's AS_PATH.	1064
authentication-key	Enables MD5 authentication and configures the authentication key.	1065
cluster	Configure the cluster ID for a route reflector server.	1066
connect-retry	Configures the BGP Connect Retry timer value in seconds.	1066
damping	Enables BGP route damping for learned routes as defined within the route policy.	1067
description	Creates a text description stored in the configuration file for a configuration context.	952
disable-client-reflect	Disable the reflection of routes by the route reflector to the clients in the given group.	1067
disable-communities	Specifies BGP to disable sending communities.	1068
disable-fast-external-failover	Configures BGP fast external failover.	1068

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
export	Specifies the export route policy to be used for determining which routes are advertised to peers.	1068
hold-time	Configures the BGP Hold Time in seconds.	1070
import	Specifies the import route policy to be used for determining which routes are accepted from peers.	1070
keepalive	Configures the BGP keepalive timer in seconds. A keepalive message is sent every time this timer expires.	1071
local-address	Configures the local IP address used by the group or neighbor when communicating with BGP peers.	1072
local-as	Adds a BGP virtual AS number.	1072
local-preference	Enables setting the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute.	1073
loop-detect	Configures how the BGP peer session will handle loop detection in the AS Path.	1073
med-out	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.	1074
min-as-origination	Configures the minimum interval in seconds at which a given path attribute, originated by the local router, can be advertised to a peer.	1075
min-route-advertisement	Configures the minimum interval in seconds at which a given prefix can be advertised to a peer.	1075
multihop	Configures the TTL value entered in the IP header of packets sent to an EBGP peer multiple hops away.	1076
next-hop-self	Configures the group or neighbor to always set the NEXTHOP path attribute to its own physical interface when advertising to an EBGP peer.	1077
passive	Enables passive mode for the BGP group or neighbor.	1077
peer-as	Configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.	1077
preference	Configures the route preference for routes learned from the configured peer(s).	1078
prefix-limit	Configures the maximum number of routes BGP can learn from a peer.	1078
remove-private	Enables removing the private AS numbers from the AS Path before advertising them to BGP peers.	1079
type	Configures the BGP peer as type internal or external.	1079
no shutdown	Administratively enables the BGP group.	951
Configure group level VPRN BGP neighbor parameters		910
config>service>vprn>bgp>group>neighbor		

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
advertise-inactive	Enables the advertising of inactive BGP routers to other BGP peers.	1063
aggregator-id-zero	Enables setting the router ID to zero in the aggregator path attribute when BGP is aggregating routes.	1063
as-override	Replaces all instances of the peer's AS number with the local AS number in a BGP route's AS_PATH.	1064
authentication-key	Enables MD5 authentication and configures the authentication key.	1065
cluster	Configure the cluster ID for a route reflector server.	1066
connect-retry	Configures the BGP Connect Retry timer value in seconds.	1066
damping	Enables BGP route damping for learned routes as defined within the route policy.	1067
description	Creates a text description stored in the configuration file for a configuration context.	952
disable-client-reflect	Disables route reflection by the route reflector to clients in the group.	1067
disable-communities	Specifies BGP to disable sending communities.	1068
disable-fast-external-failover	Configures BGP fast external failover.	1068
export	Specifies the export route policy to be used for determining which routes are advertised to peers.	1068
hold-time	Configures the BGP Hold Time in seconds.	1070
import	Specifies the import route policy to be used for determining which routes are accepted from peers.	1070
keepalive	Configures the BGP keepalive timer in seconds. A keepalive message is sent every time this timer expires.	1071
local-address	Configures the local IP address used by the group or neighbor when communicating with BGP peers.	1072
local-as	Adds a BGP virtual AS number.	1072
local-preference	Enables setting the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute.	1073
loop-detect	Configures how the BGP peer session handles loop detection in the AS Path.	1073
med-out	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.	1074
min-as-origination	Configures the minimum interval in seconds at which a given path attribute, originated by the local router, can be advertised to a peer.	1075
min-route-advertisement	Configures the minimum interval in seconds at which a given prefix can be advertised to a peer.	1075

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
multihop	Configures the TTL value entered in the IP header of packets sent to an EBGP peer multiple hops away.	1076
next-hop-self	Configures the group or neighbor to always set the NEXTHOP path attribute to its own physical interface when advertising to an EBGP peer.	1077
passive	Enables passive mode for the BGP group or neighbor.	1077
peer-as	Configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.	1077
preference	Configures the route preference for routes learned from the configured peer(s).	1078
prefix-limit	Configures the maximum number of routes BGP can learn from a peer.	1078
remove-private	Enables removing the private AS numbers from the AS Path before advertising them to BGP peers.	1079
type	Configures the BGP peer as type internal or external.	1079
no shutdown	Administratively enables the BGP neighbor.	951
Configure VPRN RIP parameters		913
config>service>vprn# rip		
authentication-key	Sets the authentication password to be passed between RIP neighbors.	1101
authentication-type	Defines the type of authentication to be used between RIP neighbors.	1102
check-zero	Enables checking for zero values in fields specified to be zero by the RIPv1 and RIPv2 specifications.	1102
description	Creates a text description for the RIP instance.	952
export	Specifies the export policies to be used to control routes advertised to RIP neighbors.	1103
import	Specifies the import policies to be used to control routes advertised from RIP neighbors.	1103
message-size	Sets the maximum number of routes per RIP update message.	1104
metric-in	Sets the metric added to routes that were received from a RIP neighbor.	1104
metric-out	Sets the metric added to routes that were exported into RIP and advertised to RIP neighbors.	1104
preference	Sets the route preference assigned to RIP routes.	1105
receive	Configures the type(s) of RIP updates that will be accepted and processed.	1105
send	Specifies the type of RIP messages sent to RIP neighbors.	1106
split-horizon	Enables the use of split-horizon.	1102
timers	Sets the values for the update, timeout, and flush timers.	1106

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
Configure VPRN RIP group parameters		912
`config>service>vprn>rip>group`		
`authentication-key`	Sets the authentication password to be passed between RIP neighbors.	1101
`authentication-type`	Defines the type of authentication to be used between RIP neighbors.	1102
`check-zero`	Enables checking for zero values in fields specified to be zero by the RIPv1 and RIPv2 specifications.	1102
`description`	Creates a text description stored in the configuration file for the RIP instance.	952
`export`	Specifies the export policies to be used to control routes advertised to RIP neighbors.	1103
`import`	Specifies the import policies to be used to control routes advertised from RIP neighbors.	1103
`message-size`	Sets the maximum number of routes per RIP update message.	1104
`metric-in`	Sets the metric added to routes that were received from a RIP neighbor.	1104
`metric-out`	Sets the metric added to routes that were exported into RIP and advertised to RIP neighbors.	1104
`preference`	Sets the route preference assigned to RIP routes.	1105
`receive`	Configures the type(s) of RIP updates that will be accepted and processed.	1105
`send`	Specifies the type of RIP messages sent to RIP neighbors.	1106
`split-horizon`	Enables the use of split-horizon.	1102
`timers`	Sets the values for the update, timeout, and flush timers.	1106
Configure VPRN RIP neighbor parameters		
`config>service>vprn>rip>group>neighbor`		
`authentication-key`	Sets the authentication password to be passed between RIP neighbors.	1101
`authentication-type`	Defines the type of authentication to be used between RIP neighbors.	1102
`check-zero`	Enables checking for zero values in fields specified to be zero by the RIPv1 and RIPv2 specifications.	1102
`description`	Creates a text description stored in the configuration file for the RIP instance.	952
`export`	Specifies the export policies to be used to control routes advertised to RIP neighbors.	1103
`import`	Specifies the import policies to be used to control routes advertised from RIP neighbors.	1103
`message-size`	Sets the maximum number of routes per RIP update message.	1104
`metric-in`	Sets the metric added to routes that were received from a RIP neighbor.	1104

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
metric-out	Sets the metric added to routes that were exported into RIP and advertised to RIP neighbors.	1104
preference	Sets the route preference assigned to RIP routes.	1105
receive	Configures the type(s) of RIP updates that will be accepted and processed.	1105
send	Specifies the type of RIP messages sent to RIP neighbors.	1106
split-horizon	Enables the use of split-horizon.	1102
timers	Sets the values for the update, timeout, and flush timers.	1106
Configure VPRN OSPF parameters		
config>service>vprn# ospf		
area	Creates the context to configure an OSPF area.	1081
area-range	Creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression.	1081
blackhole-aggregate	Installs a low priority blackhole route for the entire aggregate.	1082
interface	Creates a context to configure an OSPF interface.	1082
advertise-subnet	Enables advertising point-to-point interfaces as subnet routes (network number and mask).	1083
authentication-key	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.	1083
authentication-type	Enables authentication and specifies the type of authentication to be used on the OSPF interface.	1084
dead-interval	Configures the time, in seconds, that OSPF waits before declaring a neighbor router down.	1084
hello-interval	Configures the interval between OSPF hellos issued on the interface or virtual link	1085
interface-type	Configures the interface type to be either broadcast or point-to-point.	1085
message-digest-key	Configures a message digest key when MD5 authentication is enabled on the interface	1086
metric	Configures an explicit route cost metric for the OSPF interface that overrides the metrics calculated based on the speed of the underlying link	1087
mtu	Configures the OSPF packet size used on this interface	1087
passive	Adds the passive property to the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol.	1087

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
priority	Configures the priority of the OSPF interface that is used on election of the designated router on the subnet.	1088
retransmit-interval	Specifies the length of time, in seconds, that OSPF will wait before retransmitting an unacknowledged link state advertisement (LSA) to an OSPF neighbor.	1088
transmit-delay	Configures the estimated time, in seconds, that it takes to transmit a link state advertisement (LSA) on the interface or virtual link.	1089
nssa	Creates the context to configure an OSPF Not So Stubby Area (NSSA) and adds/removes the NSSA designation from the area.	1089
original-default-route	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).	1090
redistribute-external	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.	1090
summaries	Enables sending summary (type 3) advertisements into a stub area or Not So Stubby Area (NSSA) on an Area Border Router (ABR).	1090
stub	Enables access to the context to configure an OSPF stub area and adds/removes the stub designation from the area.	1091
default-metric	Configures the metric used by the area border router (ABR) for the default route into a stub area.	1091
virtual-link	Configures a virtual link to connect area border routers to the backbone via a virtual link.	1092
compatible-rfc1583	Enables OSPF summary and external route calculations in compliance with RFC1583 and earlier RFCs.	1093
export	Associates export route policies to determine which routes are exported from the route table to OSPF.	1093
external-db-overflow	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.	1093
external-preference	Configures the preference for OSPF external routes.	1094
overload	Changes the overload state of the local router so that it appears to be overloaded.	1095
overload-include-stub	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.	1096
overload-on-boot	Configures the IGP upon bootup in the overload state until a timeout timer expires or a manual override of the current overload state is entered with the no overload command.	1096

Table 1: CLI Commands to Configure VPRN Service Parameters (Continued)

Command	Description	Page
`preference`	Configures the preference for OSPF internal routes.	1096
`reference-bandwidth`	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.	1097
`timers`	Enables the context that allows for the configuration of OSPF timers.	1098
`spf-wait`	Defines the maximum interval between two consecutive SPF calculations.	1098
`lsa-arrival`	Defines the minimum delay that must pass between receipt of the same Link State Advertisements (LSAs) arriving from neighbors.	1099
`lsa-generate`	Customizes the throttling of OSPF LSA-generation.	1099

Basic Configuration

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

- Customer ID (refer to Configuring Customers on page 64)
- 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 ldp
            vrf-target target:10001:1
            interface "to-ce1" 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-ce1"
                    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 an route distinguisher (mandatory).
4. Define VRF route-target associations or VRF import/export policies.
5. Define PIM parameters (optional).
6. Create a subscriber interface (optional).
7. Create an interface.
8. Define SAP parameters on the interface.
   - → Select node(s) and port(s).
   - → Optional - select QoS policies other than the default (configured in `config>qos` context).
   - → Optional - select filter policies (configured in `config>filter` context).
   - → Optional - select accounting policy (configured in `config>log` context).
   - → Optional - configure DHCP features.
9. Define BGP parameters (optional).
   - → BGP must be enabled in the `config>router>bgp` context.
10. Define RIP parameters (optional).
    - → RIP must be enabled in the `config>router>rip` context.
11. Define spoke SDP parameters.
12. Enable the service.

Configuring VPRN Components

This section provides VPRN configuration examples for the following entities:

Creating a VPRN Service

Use the following CLI syntax to create a VRPN 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
```

Example:
```
config>service# vprn 1 customer 1 create
config>service>vprn# route-distinguisher 10001:0
config>service>vprn# no shutdown
```

The following example displays the VPRN service creation output.

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vprn 1 customer 1 create
            route-distinguisher 10001:0
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-1>config>service>vprn#
```

Configuring Global VPRN Parameters

Refer to VPRN Services Command Reference on page 925 for CLI syntax to configure VPRN parameters.

Example:
```
config>service# vprn 1
    config>service>vprn# vrf-import "vrfImpPolCust1"
    config>service>vprn# vrf-export "vrfExpPolCust1"
    config>service>vprn# autonomous-system 10000
    config>service>vprn# spoke-sdp 2
    config>service>vprn>sdp# no shutdown
    config>service>vprn# static-route 6.5.0.0/24 next-hop 10.1.1.2
    config>service>vprn# no shutdown
```

The following example displays the VPRN service creation output.

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vprn 1 customer 1 create
            vrf-import "vrfImpPolCust1"
            vrf-export "vrfExpPolCust1"
            autonomous-system 10000
            route-distinguisher 10001:1
            spoke-sdp 2 create
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

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|vpls}]
                egress
                   filter {ip ip-filter-id}
                   vc-label egress-vc-label
                ingress
                   filter {ip ip-filter-id}
                   vc-label ingress-vc-label
                tos-marking-state {trusted|untrusted}
                no shutdown
```

Example:
```
config>service>vprn# spoke-sdp 3 create
config>service>vprn>sdp$ no shutdown
config>service>vprn>sdp$ exit
config>service>vprn# interface "SpokeSDP" create
config>service>vprn>if$ spoke-sdp 3:4 create
config>service>vprn>if>spoke-sdp>egress$ filter ip 10
config>service>vprn>if>spoke-sdp>egress$ vc-label 2000
config>service>vprn>if>spoke-sdp>egress$ exit
config>service>vprn>if>spoke-sdp# ingress
config>service>vprn>if>spoke-sdp>ingress# filter ip 10
config>service>vprn>if>spoke-sdp>ingress# vc-label 3000
config>service>vprn>if>spoke-sdp>ingress# exit
```

```
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 VPRN Protocols - PIM

Refer to VPRN Services Command Reference on page 925 for CLI syntax to configure VPRN parameters.

The following displays the command usage to configure VPRN PIM parameters:

Example:
```
config>service# vprn 1 customer 2 create
config>service>vprn# route-distinguisher 1:11
config>service>vprn# pim
config>service>vprn>pim# interface if1
config>service>vprn>pim>if$ address 12.13.14.15/32
config>service>vprn>pim>if# loopback
config>service>vprn>pim>if# exit
config>service>vprn>pim# interface if2
config>service>vprn>pim>if$ address 14.14.14.1/24
config>service>vprn>pim>if# sap 1/1/2:0
config>service>vprn>pim>if# exit
config>service>vprn>pim# rp
config>service>vprn>pim>rp# static
config>service>vprn>pim>rp# bsr-candidate
config>service>vprn>pim>rp>bsr-cand# exit
config>service>vprn>pim>rp# rp-candidate
config>service>vprn>pim>rp>rp-cand# exit
config>service>vprn>pim>rp# exit
config>service>vprn>pim# exit
```

The following example displays the VPRN PIM configurations:

```
config>service# info
#------------------------------------------
...
        vprn 1 customer 2 create
            route-distinguisher 1:11
            interface "if1" create
                address 12.13.14.15/32
                loopback
            exit
            interface "if2" create
                address 14.14.14.1/24
                sap 1/1/2:0 create
                exit
            exit
            pim
                interface "if1"
                exit
                interface "if2"
                exit
                rp
                    static
                    exit
                    bsr-candidate
                        shutdown
                    exit
```

```
                    rp-candidate
                        shutdown
                    exit
                exit
            exit
            no shutdown
        exit
    exit
#------------------------------------------
config>service#
```

Configuring Router Interfaces

Refer to the *7750 SR OS Router Configuration Guide* for command descriptions and syntax information to configure router interfaces.

The following displays the command usage to configure interfaces used in service PIM configurations:

Example:
```
config>router# interface if1
config>router>if$ address 2.2.2.1/24
config>router>if# port 1/1/33
config>router>if# no shutdown
config>router>if# exit
config>router# interface if2
config>router>if$ address 10.49.1.46/24
config>router>if$ port 1/1/34
config>router>if$ no shutdown
config>router>if$ exit
config>router# interface if3
config>router>if$ address 11.11.11.1/24
config>router>if$ port 1/1/35
config>router>if$ no shutdown
config>router>if$ exit
config>router#
```

The following example displays the 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
...
#------------------------------------------
```

Configuring VPRN Protocols - BGP

Configuring BGP between the PE routers allows the PE routers to exchange information about routes originating and terminating in the VPRN. The PE routers use the information to determine which labels are used for traffic intended for remote sites.

In order to enable a VPRN BGP instance, the BGP protocol must be enabled in the `config>router>bgp` context on each participating SR-Series router.

NOTE: Careful planning is essential to implement commands that can affect the behavior of VPRN BGP global, group, and neighbor levels. Because the BGP commands are hierarchical, analyze the values that can disable features on a particular level.

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 904.
- Specify a router ID - Note that if a new or different router ID value is entered in the BGP context, then the new valuestakes precedence and overwrites the router-level router ID. See Configuring Global VPRN Parameters on page 904.
- Specify a VPRN BGP peer group.
- Specify a VPRN BGP neighbor with which to peer.
- Specify a VPRN BGP peer-AS that is associated with the above peer.

VPRN BGP is administratively enabled upon creation. Minimally, to enable VPRN BGP in a VPRN instance, you must associate an autonomous system number and router ID for the VPRN service, create a peer group, neighbor, and associate a peer AS number. There are no default VPRN BGP groups or neighbors. Each VPRN BGP group and neighbor must be explicitly configured.

All parameters configured for VPRN BGP are applied to the group and are inherited by each peer, but a group parameter can be overridden on a specific basis. VPRN BGP command hierarchy consists of three levels:

- The global level
- The group level
- The neighbor level

For example:

CLI Syntax:
```
config>service>vprn>bgp#                (global level)
    group                               (group level)
 neighbor                               (neighbor level)
```

For more information about the BGP protocol, refer to the *7750 SR 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.

VPRN BGP CLI Syntax

Use the CLI syntax to configure VPRN BGP parameters (BGP Configuration Commands on page 939).

The following displays the command usage to configure VPRN BGP parameters:

Example:
```
config>service>vprn# bgp
config>service>vprn>bgp# router-id 10.0.0.1
config>service>vprn>bgp# group to-cel
config>service>vprn>bgp>group# export vprnBgpExpPolCust1
config>service>vprn>bgp>group# peer-as 65101
config>service>vprn>bgp>group# neighbor 10.1.1.2
config>service>vprn>bgp>group>neighbor# exit
```

The following example displays the 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 ldp
            vrf-target target:10001:1
            interface "to-ce1" 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
            spoke-sdp 2 create
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

Configuring VPRN Protocols - RIP

PE routers which attach to a particular VPN need to know, for each of that VPN's sites, which addresses in that VPN are at each site. There are several ways that a PE router can obtain this set of addresses. The Routing Information Protocol (RIP) sends routing update messages that include entry changes. The routing table is updated to reflect the new information.

RIP can be used as a PE/CE distribution technique. PE and CE routers may be RIP peers, and the CE may use RIP to tell the PE router the set of address prefixes which are reachable at the CE router's site. When RIP is configured in the CE, care must be taken to ensure that address prefixes from other sites (i.e., address prefixes learned by the CE router from the PE router) are never advertised to the PE. Specifically, if a PE router receives a VPN-IPv4 route, and as a result distributes an IPv4 route to a CE, then that route must not be distributed back from that CE's site to a PE router (either the same router or different routers).

In order to enable a VPRN RIP instance, the RIP protocol must be enabled in the `config>router>rip` context on each participating SR-Series router. VPRN RIP is administratively enabled upon creation. Configuring other RIP commands and parameters are optional.

NOTE: Careful planning is essential to implement commands that can affect the behavior of VPRN RIP global, group, and neighbor levels. Because the RIP commands are hierarchical, analyze the values that can disable features on a particular level.

The parameters configured on the VPRN RIP global level are inherited by the group and neighbor levels. Many of the hierarchical VPRN RIP commands can be modified on different levels. The most specific value is used. That is, a VPRN RIP group-specific command takes precedence over a global VPRN RIP command. A neighbor-specific statement takes precedence over a global VPRN RIP and group-specific command. For example, if you modify a VPRN RIP neighbor-level command default, the new value takes precedence over VPRN RIP group- and global-level settings. There are no default VPRN RIP groups or neighbors. Each VPRN RIP group and neighbor must be explicitly configured.

The minimal parameters that should be configured for a VPRN instance are:

- Specify a VPRN RIP peer group.
- Specify a VPRN RIP neighbor with which to peer.
- Specify a VPRN RIP peer-AS that is associated with the above peer.

VPRN RIP command hierarchy consists of three levels:

- The global level
- The group level
- The neighbor level

For example:

CLI Syntax:
```
config>service>vprn>rip#                (global level)
    group                               (group level)
  neighbor                              (neighbor level)
```

VPRN RIP CLI Syntax

Use the CLI syntax to configure VPRN RIP parameters (RIP Configuration Commands on page 945).

The following displays the command usage to configure VPRN RIP parameters:

Example:
```
config>service>vprn# rip
config>service>vprn>rip# export vprnRipExpPolCust1
config>service>vprn>rip# group "cel"
config>service>vprn>rip>group# neighbor "to-cel"
config>service>vprn>rip>group>neighbor# exit
config>service>vprn>rip>group# exit
config>service>vprn>rip#
```

The following example displays the VPRN RIP configuration:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vprn 1 customer 1 create
            vrf-import "vrfImpPolCust1"
            vrf-export "vrfExpPolCust1"
            ecmp 8
            autonomous-system 10000
            route-distinguisher 10001:1
            auto-bind ldp
            vrf-target target:10001:1
            interface "to-ce1" 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-ce1"
                    exit
                exit
            exit
            spoke-sdp 2 create
            exit
            no shutdown
        exit
...
-------------------------------------------
*A:ALA-1>config>service# info
```

For more information about the RIP protocol, refer to the *7750 SR OS Router Configuration Guide*.

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 *7750 SR OS Router Configuration Guide*.

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

VPRN OSPF CLI Syntax

Refer to OSPF Configuration Commands on page 943 for CLI syntax to configure VPRN parameters.

The following example displays the command usage to create an VPRN OSPF interface and configure area parameters:

Example:
```
config>service# vprn 2 customer 1 create
config>service>vprn# ospf
config>service>vprn>ospf# interface test create
config>service>vprn>ospf>if$ exit
config>service>vprn>ospf>if$ no shutdown
config>service>vprn>ospf# area 0.0.0.0
config>service>vprn>ospf>area# virtual-link 1.2.3.4
          transit-area 1.2.3.4
config>service>vprn>ospf>area>virtual-link# hello-interval 9
config>service>vprn>ospf>area>virtual-link# dead-interval 40
config>service>vprn>ospf>area>virtual-link# exit
```

The following example displays the VPRN OSPF configuration shown above:

```
A:debby-sim1>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:debby-sim1>config>service#
```

For more information about the OSPF protocol, refer to the *7750 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 943 for CLI commands and syntax.

The following displays the command usage to configure VPRN interface parameters:

Example:
```
config>service# vprn 1
config>service>vprn# interface "to-ce1" create
config>service>vprn>if$ address 10.1.1.1/24
config>service>vprn>if# no shutdown
```

The following example displays the 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 ldp
            vrf-target target:10001:1
            interface "to-ce1" 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.

VPRN interface ATM SAP parameters can only be configured on ATM-type MDAs and ATM-configured ports. The `periodic-loopback` command can only be enabled when the `config>system>atm>oam` context is enabled. See the *7750 SR OS Basic System Configuration Guide.*

Refer to OSPF Configuration Commands on page 943 for CLI commands and syntax.

The following displays the command usage to configure VPRN interface SAP parameters:

Example:
```
config>service# vprn 1
config>service>vprn# interface "to-cel"
config>service>vprn>if# sap 1/1/10:1
config>service>vprn>if>sap# egress
config>service>vprn>if>sap>egress# filter ip 6
config>service>vprn>if>sap>egress# qos 1010
config>service>vprn>if>sap>egress# scheduler-policy SLA1
config>service>vprn>if>sap>egress# exit
config>service>vprn>if>sap# ingress
config>service>vprn>if>sap>ingress# qos 100
config>service>vprn>if>sap>ingress# scheduler-policy SLA2
config>service>vprn>if>sap>ingress# exit
```

The following example displays the 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 ldp
            vrf-target target:10001:1
            interface "to-ce1" 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
            spoke-sdp 2 create
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

Service Management Tasks

This section discusses the following service management tasks:

- Modifying VPRN Service Parameters on page 919
- Deleting a VPRN Service on page 921

Modifying VPRN Service Parameters

Use the CLI syntax to modify VPRN parameters (VPRN Services Command Reference on page 925).

The following displays the command usage to configure additional VPRN parameters and remove RIP from the configuration:

Example:
```
config>service# vprn 1
config>service>vprn# ecmp 8
config>service>vprn# maximum-routes 2000
config>service>vprn# rip
config>service>vprn>rip# shutdown
config>service>vprn>rip# exit
config>service>vprn# no rip
config>service>vprn#
```

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-ce1" 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-ce1"
                    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
```

Example:
```
config>service# vprn 1
config>service>vprn# interface "to-cel"
config>service>vprn>if# sap 1/1/10:1
config>service>vprn>if>sap# shutdown
config>service>vprn>if>sap# exit
config>service>vprn>if# no sap 1/1/10:1
config>service>vprn>if# shutdown
config>service>vprn>if# exit
config>service>vprn# no interface "to-cel"
config>service>vprn# bpg
config>service>vprn>bgp# shutdown
config>service>vprn>bgp# exit
config>service>vprn# no bgp
config>service>vprn# rip
config>service>vprn>rip# shutdown
config>service>vprn>rip# exit
config>service>vprn# no rip
config>service>vprn# spoke-sdp 2
config>service>vprn>sdp# shutdown
config>service>vprn>sdp# exit
config>service>vprn# shutdown
config>service>vprn# exit
config>service# no vprn 1
```

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 ldp
            vrf-target target:10001:1
            interface "to-ce1" 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-ce1"
                    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
```

Example:
```
config>service# vprn 1
config>service>vprn# no shutdown
config>service>vprn# exit
```

VPRN Services Command Reference

Command Hierarchies

VPRN Service Configuration Commands

```
config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — auto-bind {ldp | gre}
            — no auto-bind
            — autonomous-system as-number
            — no autonomous-system
            — description description-string
            — no description
            — ecmp max-ecmp-routes
            — no ecmp
            — gsmp
                — [no] group name
                    — ancp
                        — [no] dynamic-topology-discover
                        — [no] oam
                    — description description-string
                    — no description
                    — hold-multiplier multiplier
                    — no hold-multiplier
                    — keepalive seconds
```

```
                    — no keepalive
                    — [no] neighbor ip-address
                        — description description-string
                        — no description
                        — local-address ip-address
                        — no local-address
                        — priority-marking dscp dscp-name
                        — priority-marking prec ip-prec-value
                        — no priority-marking
                        — [no] shutdown
                    — [no] shutdown
                — [no] shutdown
            — igmp
                — [no] interface ip-int-name
                    — import policy-name
                    — no import
                    — max-groups value
                    — no max-groups
                    — mcac
                        — mc-constraints
                            — level level-id bw bandwidth
                            — no level level-id
                            — number-down number-lag-port-down
                            — number-down number-lag-port-down level level-id
                            — [no] shutdown
                        — policy policy-name
                        — no policy
                        — unconstrained-bw bandwidth mandatory-bw mandatory-bw
                        — no unconstrained-bw
                    — [no] shutdown
                    — 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
            — maximum-routes number [log-only] [threshold percent]
            — no maximum-routes
            — mc-maximum-routes number [log-only] [threshold percent]
            — no mc-maximum-routes
            — [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**

- **route-distinguisher** [*ip-address*:*number1* | *asn*:*number2*]
- **no route-distinguisher**
- **router-id** *ip-address*
- **no router-id**
- [**no**] **shutdown**
- **snmp-community** *community-name* [**version** *SNMP-version*]
- **no snmp-community** *community-name*
- **source-address**
- **no source-address**
  - **application** *app* [*ip-int-name*|*ip-address*]
  - **no application** *app*
- [**no**] **spoke-sdp** *sdp-id*
  - [**no**] **shutdown**
- [**no**] **static-route** {*ip-prefix*/*mask* | *ip-prefix netmask*}[**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **next-hop** *ip-address* [**bfd-enable**]
- [**no**] **static-route** {*ip-prefix*/*mask* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **indirect** *ip-address*
- [**no**] **static-route** {*ip-prefix*/*mask* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **black-hole**
- **type** *{hub | subscriber-split-horizon}*
- **no type**
- **vrf-export** *policy-name* [*policy-name*...(upto 5 max)]
- **no vrf-export**
- **vrf-import** *policy-name* [*policy-name*...(upto 5 max)]
- **no vrf-import**
- **vrf-target** {*ext-comm*|{[**export** *ext-comm*][**import** *ext-comm*]}}
- **no vrf-target**
- [**no**] **shutdown**

Subscriber Interface Commands

```
config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — subscriber-interface ip-int-name [fwd-service service-id fwd-subscriber-interface ip-int-name] [create]
            — no subscriber-interface ip-int-name
                — address {ip-address/mask | ip-address netmask} [gw-ip-address ip-address]
                — no address
                — authentication-policy name
                — no authentication-policy
                — description description-string
                — no description
                — dhcp
                    — description description-string
                    — no description
                    — gi-address ip-address [src-ip-addr]
                    — no gi-address
                    — lease-populate nbr-of-leases
                    — no lease-populate
                    — [no] option
                        — [no] vendor-specific-option
                            — [no] client-mac-address
                            — [no] sap-id
                            — [no] service-id
                            — string text
                            — no string
                            — [no] system-id
                    — proxy-server
                        — emulated-server ip-address
                        — no emulated-server
                        — lease-time [days days] [hrs hours] [min minutes] [sec seconds] [radius-override]
                        — no lease-time
                        — [no] shutdown
                    — server server1 [server2...(up to 8 max)]
                    — no server
                    — [no] shutdown
                — [no] group-interface ip-int-name
                    — [no] arp-populate
                    — arp-timeout seconds
                    — no arp-timeout
                    — authentication-policy name
                    — no authentication-policy
                    — description description-string
                    — no description
                    — dhcp
                        — description description-string
                        — no description
                        — gi-address ip-address [src-ip-addr]
                        — no gi-address
                        — lease-populate nbr-of-leases
                        — no lease-populate
                        — [no] match-circuit-id
```

- [no] option
  - action {replace | drop | keep}
  - no action
  - circuit-id [ascii-tuple|ifindex|sap-id|vlan-ascii-tupl]
  - no circuit-id
  - remote-id [mac | string *string*]
  - no remote-id
  - [no] vendor-specific-option
    - [no] client-mac-address
    - [no] sap-id
    - [no] service-id
    - string *text*
    - no string
    - [no] system-id
  - proxy-server
    - emulated-server *ip-address*
    - no emulated-server
    - lease-time [days *days*] [hrs *hours*] [min *minutes*] [sec *seconds*] [radius-override]
    - no lease-time
    - [no] shutdown
  - server *server1* [*server2*...(up to 8 max)]
  - no server
  - [no] shutdown
  - [no] trusted
- host-connectivity-verify [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
- [no] local-proxy-arp
- [no] mac *ieee-address*
- [no] proxy-arp-policy *policy-name* [*policy-name*...(up to 5 max)]
- redundant-interface *red-ip-int-name*
- no redundant-interface
- [no] remote-proxy-arp
- [no] sap *sap-id*
  - accounting-policy *acct-policy-id*
  - no accounting-policy [*acct-policy-id*]
  - anti-spoof {ip | ip-mac}
  - no anti-spoof
  - atm
    - egress
      - traffic-desc *traffic-desc-profile-id*
      - no traffic-desc
    - encapsulation *atm-encap-type*
    - ingress
      - traffic-desc *traffic-desc-profile-id*

```
                — no traffic-desc
            — oam
                — [no] alarm-cells
                — [no] periodic-loopback
        — [no] collect-stats
        — description description-string
        — no description
        — egress
            — agg-rate-limit agg-rate
            — no agg-rate-limit agg-rate
            — filter ip ip-filter-id
            — filter ipv6 ipv6-filter-id
            — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
            — no filter
            — [no] qinq-mark-top-only
            — [no] qos policy-id
            — scheduler-policy scheduler-policy-name
            — no scheduler-policy
        — host ip ip-address [mac ieee-address] [subscriber sub-ident-string] [sub-profile sub-profile-name] [sla-profile sla-profile-name] [ancp-string ancp-string]
        — no host {[ip ip-address] [mac ieee-address]}
        — no host all
        — ingress
            — filter ip ip-filter-id
            — filter ipv6 ipv6-filter-id
            — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
            — no filter
            — match-qinq-dot1p {top|bottom}
            — no match-qinq-dot1p
            — qos policy-id [shared-queuing]
            — no qos policy-id
            — scheduler-policy scheduler-policy-name
            — no scheduler-policy
        — multi-service-site customer-site-name
        — no multi-service-site
        — [no] shutdown
        — [no] sub-sla-mgmt
            — def-sla-profile default-sla-profile-name
            — no def-sla-profile
            — def-sub-profile default-subscriber-profile-name
            — no def-sub-profile
            — multi-sub-sap subscriber-limit
            — no multi-sub-sap
            — [no] shutdown
            — single-sub-parameters
                — non-sub-traffic sub-profile sub-profile-name sla-profile sla-profile-name [subscriber sub-ident-string]
                — no non-sub-traffic
                — [no] profiled-traffic-only
            — sub-ident-policy sub-ident-policy-name
            — no sub-ident-policy
    — [no] shutdown
    — [no] srrp srrp-id
            — description description-string
            — no description
```

```
                    — gw-mac mac-address
                    — no gw-mac
                    — keep-alive-interval interval
                    — no keep-alive-interval
                    — message-path sap-id
                    — no message-path
                    — [no] policy vrrp-policy-id
                    — priority priority
                    — no priority
                    — [no] shutdown
                — tos-marking-state {trusted | untrusted}
                — no tos-marking-state
            — [no] shutdown
```

Interface Commands

```
config
    — service
        — vprn
            — [no] interface ip-int-name
                — [no] active-cpm-protocols
                — address ip-address[/mask] [netmask] [broadcast {all-ones | host-ones}]
                — no address
                — [no] allow-directed-broadcasts
                — [no] arp-populate
                — arp-timeout [seconds]
                — no arp-timeout
                — authentication-policy name
                — no authentication-policy
                — bfd transmit-interval [receive receive-interval] [multiplier multiplier]
                — no bfd
                — cflowd {acl | interface}
                — no cflowd
                — description description-string
                — no description [description-string]
                — dhcp
                    — description description-string
                    — no description
                    — gi-address ip-address [src-ip-addr]
                    — no gi-address
                    — lease-populate [nbr-of-leases]
                    — no lease-populate
                    — [no] option
                        — action {replace | drop | keep}
                        — no action
                        — circuit-id [ascii-tuple | ifindex | sap-id |vlan-ascii-tuple]
                        — no circuit-id
                        — remote-id  [mac | string string]
                        — [no] vendor-specific-option
                            — [no] client-mac-address
                            — [no] sap-id
                            — [no] service-id
                            — string text
                            — no string
                            — [no] system-id
                    — proxy-server
                        — emulated-server ip-address
                        — no emulated-server
                        — lease-time [days days] [hrs hours] [min minutes] [sec seconds] [radius-override]
                        — no lease-time
                        — [no] shutdown
                    — server server1 [server2...(up to 8 max)]
                    — no server
                    — [no] shutdown
                    — [no] trusted
                — 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*]

- **ip-mtu** *octets*
- **no ip-mtu**
- [**no**] **local-proxy-arp**
- [**no**] **loopback**
- **mac** *ieee-address*
- **no mac** [*ieee-address*]
- [**no**] **proxy-arp-policy**
- [**no**] **remote-proxy-arp**
- **secondary** {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**] [**igp-inhibit**]
- **no secondary** {*ip-address/mask* | *ip-address netmask*}
- [**no**] **shutdown**
- **spoke-sdp** *sdp-id*[:*vc-id***]**
- **no spoke-sdp** *sdp-id*[:*vc-id*]
  - **accounting-policy** *acct-policy-id*
  - **no accounting-policy**
  - [**no**] **collect-stats**
  - **egress**
    - **filter** {**ip** *ip-filter-id*}
    - **no filter**
    - **vc-label** ***egress-vc-label***
    - **no vc-label** [***egress-vc-label***]
  - **ingress**
    - **filter** {**ip** *ip-filter-id*}
    - **no filter**
    - **vc-label** *ingress-vc-label*
    - **no vc-label** [*ingress-vc-label*]
  - [**no**] **shutdown**
- **static-arp** *ip-address ieee-address*
- **no static-arp** *ip-address* [*ieee-address*]
- **tos-marking-state** {**trusted** | **untrusted**}
- **no tos-marking-state**
- **unnumbered** [*ip-int-name*| *ip-address*]
- **no unnumbered**

Interface SAP Commands

```
config
    — service
        — vprn
            — [no] interface ip-int-name
                — [no] sap sap-id
                    — accounting-policy acct-policy-id
                    — no accounting-policy [acct-policy-id]
                    — anti-spoof {ip | mac | ip-mac}
                    — no anti-spoof
                    — atm
                        — egress
                            — traffic-desc traffic-desc-profile-id
                            — no traffic-desc
                        — encapsulation atm-encap-type
                        — ingress
                            — traffic-desc traffic-desc-profile-id
                            — no traffic-desc
                        — oam
                            — [no] alarm-cells
                            — [no] periodic-loopback
                    — [no] collect-stats
                    — description description-string
                    — no description [description-string]
                    — egress
                        — agg-rate-limit agg-rate
                        — no agg-rate-limit agg-rate
                        — filter ip ip-filter-id
                        — no filter [ip ip-filter-id]
                        — [no] qinq-mark-top-only
                        — qos policy-id
                        — no qos [policy-id]
                        — [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
                                — rate pir-rate [cir cir-rate]
                                — no rate
                        — scheduler-policy scheduler-policy-name
                        — no scheduler-policy
                    — host {[ip ip-address] [mac ieee-address]} [subscriber sub-ident-string] [sub-profile sub-profile-name] [sla-profile sla-profile-name]
```

- **no host** {[**ip** *ip-address*] [**mac** *ieee-address*]} [**all**]
- **ingress**
  - **filter ip** *ip-filter-id*
  - **no filter** [**ip** *ip-filter-id*]
  - **match-qinq-dot1p** {**top** | **bottom**}
  - **qos** ***policy-id*** [**shared-queuing** | **multipoint-shared**]
  - **no qos** [***policy-id***]
  - [**no**] **queue-override**
    - [**no**] **queue** *queue-id*
      - **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]
      - **no adaptation-rule**
      - **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*
      - **rate** *pir-rate* [**cir** *cir-rate*]
      - **no rate**
  - **scheduler-policy** *scheduler-policy-name*
  - **no scheduler-policy**
- **multi-service-site** *customer-site-name*
- **no multi-service-site**
- [**no**] **shutdown**
- **tod-suite** ***tod-suite-name***
- **no tod-suite**

Interface VRRP Commands

```
config
    — service
        — vprn
            — interface ip-int-name
                — vrrp virtual-router-id [owner]
                — no vrrp virtual-router-id
                    — authentication-key {authentication-key | hash-key} [hash | hash2]
                    — no authentication-key
                    — authentication-type {password | message-digest}
                    — no authentication-type
                    — [no] backup 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
```

PIM Configuration Commands

```
config
    — service
        — vprn
            — [no] pim
                — apply-to {all | none}
                — import {join-policy | register-policy} [policy-name [.. policy-name]
                — no import {join-policy | register-policy}
                — [no] interface ip-int-name
                    — [no] bfd-enable
                    — [no] bsm-check-rtr-alert
                    — hello-interval hello-interval
                    — no hello-interval
                    — hello-multiplier deci-units
                    — no hello-multiplier
                    — [no] improved-assert
                    — max-groups value
                    — no max-groups
                    — mcac
                        — mc-constraints
                            — level level-id bw bandwidth
                            — no level
                            — number-down number-lag-port-down
                            — no number-down
                            — [no] shutdown
                        — policy policy-name
                        — no policy
                        — unconstrained-bw bandwidth mandatory-bw mandatory-bw
                        — no unconstrained-bw
                    — multicast-senders {auto | always | never}
                    — no multicast-senders
                    — priority dr-priority
                    — no priority
                    — [no] shutdown
                    — sticky-dr [priority dr-priority]
                    — no sticky-dr
                    — three-way-hello [compatibility-mode]
                    — no three-way-hello
                    — [no] tracking-support
                — mdt
                    — data {grp-ip-address/mask | grp-ip-address netmask}
                    — no data
                    — data-delay-interval value
                    — no data-delay-interval
                    — data-threshold {c-grp-ip-address/mask | c-grp-ip-address netmask} mdt-threshold
                    — no data-threshold {c-grp-ip-address/mask | c-grp-ip-address netmask}
                    — default grp-ip-address
                    — no default
                        — hello-interval hello-interval
                        — no hello-interval
                        — hello-multiplier deci-units
```

```
                    — no hello-multiplier
                    — [no] improved-assert
                    — [no] shutdown
                    — three-way-hello [compatibility-mode]
                    — no three-way-hello
                    — [no] tracking-support
               — [no] join-tlv-packing-disable
          — [no] non-dr-attract-traffic
          — rp
               — [no] anycast rp-ip-address
                    — [no] rp-set-peer ip-address
               — bootstrap-export policy-name [.. policy-name...up to 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
               — rp-candidate
                    — address ip-address
                    — no address
                    — [no] group-range {grp-ip-address/mask | grp-ip-
                      address [netmask]}
                    — holdtime holdtime
                    — no holdtime
                    — priority priority
                    — no priority
                    — [no] shutdown
               — static
                    — [no] address ip-address
                      — [no] group-prefix {grp-ip-address/mask | grp-ip-
                        address netmask}
                      — [no] override
          — [no] shutdown
          — spt-switchover-threshold {grp-ip-address/mask | grp-ipaddress netmask}
            spt-threshold
          — no spt-switchover-threshold {grp-ip-address/mask | grp-ipaddress net-
            mask}
          — [no] ssm-groups
               — [no] group-range {grp-ip-address/mask | grp-ip-address net-
                 mask}
```

BGP Configuration Commands

```
config
    — service
        — vprn
            — [no] bgp
                — [no] advertise-inactive
                — [no] aggregator-id-zero
                — always-compare-med {zero | infinity}
                — no always-compare-med
                — [no] as-path-ignore
                — auth-keychain name
                — authentication-key [authentication-key | hash-key] [hash | hash2]
                — no authentication-key
                — cluster cluster-id
                — no cluster
                — [no] connect-retry seconds
                — [no] damping
                — description description-string
                — no description
                — [no] disable-client-reflect
                — disable-communities [standard] [extended]
                — no disable-communities
                — [no] disable-fast-external-failover
                — [no] enable-peer-tracking
                — export policy-name [policy-name...(upto 5 max)]
                — no export
                — hold-time seconds
                — 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-as-origination seconds
                — no min-as-origination
                — min-route-advertisement seconds
                — no min-route-advertisement
                — multihop ttl-value
                — no multihop
                — multipath max-paths
                — no multipath
                — preference preference
                — no preference
                — [no] remove-private
                — router-id ip-address
                — no router-id
```

— [**no**] **shutdown**
— [**no**] **group** *name*
— [**no**] **advertise-inactive**
— [**no**] **aggregator-id-zero**
— [**no**] **as-override**
— **auth-keychain** *name*
— **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
— **no authentication-key**
— **cluster** *cluster-id*
— **no cluster**
— **connect-retry** *seconds*
— **no connect-retry**
— [**no**] **damping**
— **description** *description-string*
— **no description**
— [**no**] **disable-client-reflect**
— **disable-communities** [**standard**] [**extended**]
— **no disable-communities**
— [**no**] **disable-fast-external-failover**
— [**no**] **enable-peer-tracking**
— **export** *policy-name* [*policy-name*...(upto 5 max)]
— **no export**
— **hold-time** *seconds*
— **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 local-as**
— **local-preference** *local-preference*
— **no local-preference**
— **loop-detect** {**drop-peer**|**discard-route**|**ignore-loop**|**off**}
— **no loop-detect**
— **med-out** {**number** | **igp-cost**}
— **no med-out**
— **min-as-origination** *seconds*
— **no min-as-origination**
— **min-route-advertisement** *seconds*
— **no min-route-advertisement**
— **multihop** *ttl-value*
— **no multihop**
— [**no**] **next-hop-self**
— [**no**] **passive**
— **peer-as** *as-number*
— **no peer-as**
— **preference** *preference*
— **no preference**
— **prefix-limit** *limit*
— **no prefix-limit**
— [**no**] **remove-private**
— [**no**] **shutdown**
— **ttl-security** *min-ttl-value*
— **no ttl-security**

- **type** {**internal** | **external**}
- **no type**
- [**no**] **neighbor** *ip-address*
  - [**no**] **advertise-inactive**
  - [**no**] **aggregator-id-zero**
  - [**no**] **as-override**
  - **auth-keychain** *name*
  - **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
  - **no authentication-key**
  - **cluster** ***cluster-id***
  - **no cluster**
  - **connect-retry** *seconds*
  - **no connect-retry**
  - [**no**] **damping**
  - **description** *description-string*
  - **no description**
  - [**no**] **disable-client-reflect**
  - **disable-communities** [**standard**] [**extended**]
  - **no disable-communities**
  - [**no**] **disable-fast-external-failover**
  - [**no**] **enable-peer-tracking**
  - **export** *policy-name* [*policy-name*...(upto 5 max)]
  - **no export**
  - **hold-time** *seconds*
  - **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 local-as**
  - **local-preference** *local-preference*
  - **no local-preference**
  - **loop-detect** {**drop-peer**|**discard-route**|**ignore-loop**|**off**}
  - **no loop-detect**
  - **med-out** {**number** | **igp-cost**}
  - **no med-out**
  - **min-as-origination** *seconds*
  - **no min-as-origination**
  - **min-route-advertisement** *seconds*
  - **no min-route-advertisement**
  - **multihop** *ttl-value*
  - **no multihop**
  - [**no**] **next-hop-self**
  - [**no**] **passive**
  - **peer-as** *as-number*
  - **no peer-as**
  - **preference** *preference*
  - **no preference**
  - **prefix-limit** *limit*
  - **no prefix-limit**

— [**no**] **remove-private**
— [**no**] **shutdown**
— **ttl-security** *min-ttl-value*
— **no ttl-security**
— **type** {**internal** | **external**}
— **no type**

OSPF Configuration Commands

```
config
    — service
        — vprn
            — [no] ospf
                — [no] area area-id
                    — area-range ip-prefix/mask [advertise | not-advertise]
                    — no area-range ip-prefix/mask
                    — [no] blackhole-aggregate
                    — [no] interface ip-int-name
                        — [no] advertise-subnet
                        — 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
                        — interface-type {broadcast | point-to-point}
                        — no interface-type
                        — message-digest-key key-id md5 [key | hash-key] [hash | hash2]
                        — no message-digest-key key-id
                        — metric metric
                        — no metric
                        — mtu bytes
                        — no mtu
                        — [no] passive
                        — priority number
                        — no priority
                        — retransmit-interval seconds
                        — no retransmit-interval
                        — [no] shutdown
                        — transit-delay seconds
                        — no transit-delay
                — [no] nssa
                    — area-range ip-prefix/mask [advertise | not-advertise]
                    — no area-range ip-prefix/mask
                    — originate-default-route [type-7]
                    — no originate-default-route
                    — [no] redistribute-external
                    — [no] summaries
                — [no] 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**
- **overload** [**timeout** *seconds*]
- **no overload**
- [**no**] **overload-include-stub**
- **overload-on-boot** [**timeout** *seconds*]
- **no overload-on-boot**
- **preference** *preference*
- **no preference**
- **reference-bandwidth** *bandwidth-in-kbps*
- **no reference-bandwidth**
- **router-id** *ip-address*
- **no router-id**
- [**no**] **shutdown**
- **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*]]

RIP Configuration Commands

```
config
    — service
        — vprn
            — [no] rip
                — authentication-key [authentication-key | hash-key] [hash | hash2]
                — no authentication-key
                — authentication-type {none | password | message-digest}
                — no authentication-type
                — check-zero {enable | disable}
                — no check-zero
                — description description-string
                — no description
                — export policy-name [policy-name...(upto 5 max)]
                — no export
                — [no] group name
                    — authentication-key [authentication-key | hash-key] [hash | hash2]
                    — no authentication-key
                    — authentication-type {none | password | message-digest}
                    — no authentication-type
                    — check-zero {enable | disable}
                    — no check-zero
                    — description description-string
                    — no description
                    — export policy-name [policy-name...(upto 5 max)]
                    — no export
                    — import policy-name [policy-name...(upto 5 max)]
                    — no import
                    — message-size max-num-of-routes
                    — no message-size
                    — metric-in metric
                    — no metric-in
                    — metric-out metric
                    — no metric-out
                    — preference preference
                    — no preference
                    — receive receive-type
                    — no receive
                    — send send-type
                    — no send
                    — [no] shutdown
                    — split-horizon {enable | disable}
                    — no split-horizon
                    — timers update timeout flush
                    — no timers
                    — [no] neighbor ip-int-name
                        — authentication-key authentication-key | hash-key [hash | hash2]
                        — no authentication-key
                        — authentication-type {none | password | message-digest}
                        — no authentication-type
                        — check-zero {enable | disable}
```

- **no check-zero**
- **description** *description-string*
- **no description**
- **export** *policy-name* [*policy-name*...(upto 5 max)]
- **no export**
- **import** *policy-name* [*policy-name*...(upto 5 max)]
- **no import**
- **message-size** *max-num-of-routes*
- **no message-size**
- **metric-in** *metric*
- **no metric-in**
- **metric-out** *metric*
- **no metric-out**
- **preference** *preference*
- **no preference**
- **receive** *receive-type*
- **no receive**
- **send** *send-type*
- **no send**
- [**no**] **shutdown**
- **split-horizon** {**enable** | **disable**}
- **no split-horizon**
- **no timers**
- **timers** *update timeout flush*

- **import** *policy-name* [*policy-name*...(upto 5 max)]
- **no import**
- **message-size** *max-num-of-routes*
- **no message-size**
- **metric-in** *metric*
- **no metric-in**
- **metric-out** *metric*
- **no metric-out**
- **preference** *preference*
- **no preference**
- **receive** *receive-type*
- **no receive**
- **send** *send-type*
- **no send**
- [**no**] **shutdown**
- **split-horizon** {**enable** | **disable**}
- **no split-horizon**
- **timers** *update timeout flush*
- **no timers**

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*]
    - **base**
    - **dhcp**
      - **lease-state** [[**sap** *sap-id*] [**sdp** [*sdp-id[:vc-id]*]]] | [**interface** *interface-name*] | [**ip-address** *ip-address[/mask]*] | [**mac** *ieee-address*] | [**wholesaler** *service-id*]] [**detail**]
      - **statistics** [**sap** *sap-id*
      - **statistics** [**sdp** *sdp-id*:*vc-id*]
      - **statistics** [**interface** *interface-name*]
      - **summary**
    - **gsmp**
      - **neighbors group** [*name*] [*ip-address*]
      - **sessions** [**group** *name*] **neighbor** *ip-address*] [**port** *port-number*] [**association**] [**statistics**]
    - **host** [**sap** *sap-id*] [**detail**]
    - **host summary**
    - **host** [**detail**] **wholesaler** *service-id*
    - **interface** [*ip-address* | *ip-int-name*] [**detail**]
    - **retailers**
    - **sap** [*sap-id* [**detail**]]
    - **sdp** [*sdp-id* | **far-end** *ip-address*] [**detail**]
    - **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-address*[/*mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**detail**]
    - **subscriber-hosts** [**detail**] **wholesaler** *service-id*
    - **wholesalers**
  - **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*
  - **sdp** [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]
  - **sdp-using** [*sdp-id*[:*vc-id*]
  - **service-using** [**vprn**] [**sdp** *sdp-id*] [**customer** *customer-id*]

show
- **router** [*vprn-service-id*]
  - **aggregate** [**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]
                — dhcp
                    — statistics [ip-int-name | ip-address]
                    — summary
                — ecmp
                — interface [{[ip-address | ip-int-name] [detail]} | summary | exclude-services]
                — rip
                    — database [ip-address[/mask] [longer]] [peer ip-address] [detail]
                    — neighbor [ip-int-name | ip-address] [detail] [advertised-routes]
                    — peer [interface-name]
                    — statistics [ip-int-name | ip-address]
                — 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]
                — tunnel-table [ip-address[/mask] [protocol protocol | sdp sdp-id]
                — tunnel-table [summary]
```

Clear Commands

```
clear
    — router
        — arp {all | ip-address}
        — arp interface [ip-int-name | ip-address]
        — 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]
        — forwarding-table [slot-number]
        — interface [ip-int-name | ip-address] [icmp]
        — rip
            — database
            — statistics [neighbor ip-int-name | ip-address]
clear
    — service
        — id service-id
            — fdb {all | mac ieee-address | sap sap-id | mesh-sdp sdp-id[:vc-id] | spoke-sdp sdp-
              id:vc-id}
            — 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**}
— **stp**

Debug Commands

debug
— **service**
— **id** *service-id*
— [**no**] **dhcp**
— **detail-level** {**low** | **medium** | **high**}
— **no detail-level**
— **mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}
— **no mode**
— [**no**] **sap** *sap-id*
— [**no**] **sdp** *sdp-id:vc-id*
— [**no**] **event-type** {**config-change** | **svc-oper-status-change** | **sap-oper-status-change** | **sdpbind-oper-status-change**}
— [**no**] **sap** *sap-id*
— **event-type** {**config-change** | **oper-status-change**}
— [**no**] **sdp** *sdp-id:vc-id*
— **event-type** {**config-change** | **oper-status-change**}
— **stp**
— [**no**] **all-events**
— [**no**] **bpdu**
— [**no**] **core-connectivity**
— [**no**] **exception**
— [**no**] **fsm-state-changes**
— [**no**] **fsm-timers**
— [**no**] **port-role**
— [**no**] **port-state**
— [**no**] **sap** *sap-id*
— [**no**] **sdp** *sdp-id:vc-id*

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>ies>sub-if>grp-if
config>service>ies>sub-if>grp-if>srrp
config>service>vprn
config>service>vprn>red-if
config>service>vprn>gsmp
config>service>vprn>gsmp>group
config>service>vprn>gsmp>group>neighbor
config>service>vprn>igmp
config>service>vprn>igmp>interface
config>service>vprn>igmp>if>mcac
config>service>vprn>igmp>if>mcac>mc-constraints
config>service>vprn>interface
config>service>vprn>if>dhcp
config>service>vprn>if>dhcp>proxy
config>service>vprn>if>vrrp
config>service>vprn>if>sap
config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
config>service>vprn>ospf
config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-link
config>service>vprn>red-if>spoke-sdp
config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor
config>service>vprn>pim
config>service>vprn>pim>interface
config>service>vprn>pim>rp>bsr-candidate
config>service>vprn>pim>rp>rp-candidate
config>service>vprn>spoke-sdp
config>service>vprn>sub-if>grp-if
config>service>vprn>sub-if>grp-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp>proxy-server
config>service>vprn>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt

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>if>dhcp
config>service>vprn>bgp
config>service>vprn>rip
config>service>vprn
config>service>ies>sub-if>grp-if>srrp
config>service>vprn>red-if
config>service>vprn>interface
config>service>vprn>if>sap
config>service>vprn>if>dhcp
config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor
config>service>vprn>subscriber-interface
config>service>vprn>sub-if>dhcp
config>service>vprn>sub-if>grp-if
config>service>vprn>sub-if>grp-if>dhcp
config>service>vprn>sub-if>grp-if>sap>atm

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

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

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

Default No description 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** *service-id* [**customer** *customer-id*]
no vprn *service-id*

Context config>service

Description This command creates or edits a Virtual Private Routed Network (VPRN) service instance.

If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

VPRN services allow the creation of customer-facing IP interfaces in the same routing instance used for service network core routing connectivity. VPRN services require that the IP addressing scheme used by the subscriber must be unique between it and other addressing schemes used by the provider and potentially the entire Internet.

IP interfaces defined within the context of an VPRN service ID must have a SAP created as the access point to the subscriber network.

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

When a service is created, the use of the **customer** *customer-id* is optional to navigate into the service configuration context. If attempting to edit a service with the incorrect *customer-id* results in an error.

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

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

Default None — No VPRN service instances exist until they are explicitly created.

Parameters *service-id* — The unique service identification number identifies the service in the service domain. The ID must be unique to this service and cannot be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR on which this service is defined.

Values 1 — 2147483647

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

Values 1 — 2147483647 — The customer identifcation number must already have been created.

auto-bind

Syntax **auto-bind** {**ldp | gre**}
no auto-bind

Context config>service>vprn

Description This command specifies the automatic binding type for the SDP assigned to this service.

Default None — The auto-bind type must be explicitly specified.

Parameters **ldp —** Specifies LDP to be the automatic binding for the SDP assigned to the service.

gre — Specifies GRE to be the automatic binding for the SDP assigned to the service.

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).

The **no** form of the command removes the defined AS from this VPRN context.

Default **no autonomous-system**

Parameters *as-number* — The 16-bit AS number.

Values 1 — 65535

ecmp

Syntax **ecmp** *max-ecmp-routes*
no ecmp

Context config>service>vprn

Description This command enables equal-cost multipath (ECMP) and configures the number of routes for path sharing. For example, the value of 2 means that 2 equal cost routes will be used for cost sharing.

ECMP groups form when the system routes to the same destination with equal cost values. Routing table entries can be entered manually (as static routes), or they can be formed when neighbors are discovered and routing table informationis exchanged by routing protocols. The system can balance traffic across the groups with equal costs.

ECMP can only be used for routes learned with the same preference and same protocol. See the discussion on preferences in the **static-route** command.

When more ECMP routes are available at the best preference than configured by the max-ecmp-routes parameter, then the lowest next-hop IP address algorithm is used to select the number of routes configured.

The **no** form of the command disables ECMP path sharing. If ECMP is disabled and multiple routes are available at the best preference and equal cost, then the route with the lowest next-hop IP address is used.

Default **no ecmp**

Parameters *max-ecmp-routes* — Integer.

Values 1 — 16

gsmp

Syntax **gsmp**

Context config>service>vprn

Description This command enables the context to configure GSMP connections maintained in this service.

Default **not enabled**

group

Syntax [**no**] **group** *name*

Context config>service>vprn>gsmp

Description This command specifies a GSMP name. A GSMP group name is unique only within the scope of the service in which it is defined.

ancp

Syntax **ancp**

Context config>service>vprn>gsmp>group

Description This command configures ANCP parameters for this GSMP group.

dynamic-topology-discover

Syntax [**no**] **dynamic-topology-discover**

Context config>service>vprn>gsmp>group>ancp

Description This command enables the ANCP dynamic topology discovery capability.

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

oam

Syntax [**no**] **oam**

Context config>service>vprn>gsmp>group>ancp

Description This command specifies whether or not the GSMP ANCP OAM capability should be negotiated at startup of the GSMP connection.

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

hold-multiplier

Syntax **hold-multiplier** *multiplier*
no hold-multiplier

Context config>service>vprn>gsmp>group

Description This command configures the hold-multiplier for the GSMP connections in this group.

Parameters *multiplier* — Specifies the GSMP hold multiplier value.

Values 1 — 100

keepalive

Syntax **keepalive** *seconds*
no keepalive

Context config>service>vprn>gsmp>group

Description This command configures keepalive values for the GSMP connections in this group.

Parameters *seconds* — Specifies the GSMP keepalive timer value in seconds.

Values 1 — 25

neighbor

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

Context config>service>vprn>gsmp>group

Description This command adds or removes a neighbor in this group.

Parameters *ip-address* — Specify the IP address in dotted decimal notation.

local-address

Syntax **local-address** *ip-address*
no local-address

Context config>service>vprn>gsmp>group>neighbor

Description This command configures the source ip-address used in the connection towards the neighbor.

Parameters *ip-address* — Specify the IP address in dotted decimal notation.

priority-marking

Syntax **priority-marking dscp** *dscp-name*
priority-marking prec *ip-prec-value*
no priority-marking

Context config>service>vprn>gsmp>group>neighbor

Description This command configures the type of priority marking to be used.

Parameters **dscp** *dscp-name* — Specifies the DSCP code-point to be used.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

prec *ip-prec-value* — Specifies the precedence value to be used.

Values 0 — 7

igmp

Syntax **igmp**

Context config>service>vprn

Description This command enables the context to configure IGMP connections maintained in this service.

Default **not enabled**

interface

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

Context config>service>vprn>igmp

Description This command enables the context to configure IGMP connections maintained in this service.

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>interface

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>interface

Description This command specifies the maximum number of groups for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed.

Default 0, no limit to the number of groups.

Parameters *value* — Specifies the maximum number of groups for this interface.

Values 1 — 16000

mcac

Syntax **mcac**

Context config>service>vprn>interface
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 levels and their associated bandwidth for multicast cac policy on this interface.

Parameters *level-id* **—** Specifies has an entry for each 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** *number-lag-port-down*
no number-down

Context config>service>vprn>igmp>if>mcac>mc-constraints
config>service>vprn>pim>if>mcac>mc-constraints

Description This command configure the number of ports down along with level for multicast cac policy on this interface.

Default **not enabled**

policy

Syntax **policy** *policy-name*
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 *policy-name* — 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** *bandwidth* **mandatory-bw** *mandatory-bw*
no unconstrained-bw

Context config>service>vprn>igmp>if>mcac
config>service>vprn>pim>if>mcac

Description This command configure the bandwidth assigned for interface's multicast CAC policy traffic. When disabled (**no unconstrained-bw)**, there is no constraint on bandwidth allocated at the interface. When enabled and if a multicast CAC policy is assigned on the interface, then no group (channel) from that policy is allowed on the interface.

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 *bandwidth* — The bandwidth assigned for interface's MCAC policy traffic, in kilo-bits per second (kbps).

Values 1— 2147483647

mandatory-bw *mandatory-bw* — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilo-bits per second (kbps).

If the *bandwidth* value is 0, no mandatory channels are allowed. If the value of *bandwidth* is '-1', then all mandatory and optional channels are allowed.

If the value of *mandatory-bw* is equal to the value of *bandwidth*, 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 *mandatory-bw* should always be less than or equal to that of *bandwidth*, An attempt to set the value of *mandatory-bw* greater than that of *bandwidth*, will result in inconsistent value error.

Values 1— 2147483647

static

Syntax **static**

Context config>service>vprn>igmp>interface

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** *grp-ip-address*

Context config>service>vprn>igmp>interface>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. Specify the address in dotted decimal notation

source

Syntax **source**

Context config>service>vprn>igmp>interface>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 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>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**

subnet-check

Syntax [**no**] **subnet-check**

Context config>service>vprn>igmp>interface

Description This command enables subnet checking for IGMP messages received on this interface. All IGMP packets with a source address that is not in the local subnet are dropped.

Default **enabled**

version

Syntax **version** *version*
no version

Context config>service>vprn>igmp>interface

Description This command specifies the IGMP version. If routers run different versions of IGMP, they will negotiate the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version. For IGMP to function correctly, all routers on a LAN should be configured to run the same version of IGMP on that LAN.

For IGMPv3, 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.

Default **3**

Parameters *version* — Specifies the IGMP version number.

Values 1, 2, 3

query-interval

Syntax **query-interval** *seconds*
no query-interval

Context config>service>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 *seconds* — 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** *seconds*

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 *seconds* — Specifies the frequency, in seconds, at which query messages are sent.

Values 1 — 1024

query-response-interval

Syntax **query-response-interval** *seconds*

Context config>service>vprn>igmp

Description This command specifies how long the querier router waits to receive a response to a host-query message from a host.

Default **10**

Parameters *seconds* — Specifies the the length of time to wait to receive a response to the host-query message from the host.

Values 1 — 1023

robust-count

Syntax **robust-count** *robust-count*
no robust-count

Context config>service>vprn>igmp

Description This command configures the robust count. The robust-count 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 *robust-count* — Specifies the robust count value.

Values 2 — 10

ssm-translate

Syntax **igmp**

Context config>service>vprn>igmp

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>igmp>ssm-translate

Description This command is used to configure group ranges which are translated to SSM (S,G) entries.

Parameters *start* — An IP address that specifies the start of the group range.

end — An IP address that specifies the end of the group range. This value should always be greater than or equal to the value of the *start* value.

source

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

Context config>service>vprn>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.

maximum-routes

Syntax **maximum-routes** *number* [**log-only**] [**threshold** *percentage*]
no maximum-routes

Context config>service>vprn

Description This command specifies the maximum number of routes that can be held within a VPN routing/forwarding (VRF) context.

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 VPRN service ID must be in a shutdown state in order to modify maximum-routes command parameters.

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 **no maximum-routes**

Parameters *number* — 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 *percentage* — The percentage at which a warning log message and SNMP trap should be sent.

Values 1 — 99

Default 10

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

route-distinguisher

Syntax **route-distinguisher** [*ip-address*:*number* | *asn*:*number*]
no route-distinguisher

Context config>service>vprn

Description This command sets the identifier attached to routes the VPN belongs to. Each routing instance must have a unique (within the carrier's domain) route distinguisher associated with it. A route distinguisher must be defined for a VPRN to be operationally active.

Default **no route-distinguisher**

Parameters The route distinguisher is a 6-byte value that can be specified in one of the following formats:

ip-address:number — Specify the IP address in dotted decimal notation. The assigned number must not be greater than 65535.

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.

router-id

Syntax **router-id** *ip-address*
no router-id

Context config>service>vprn
config>service>vprn>ospf
config>service>vprn>bgp

Description This command sets the router ID for a specific VPRN context.

If neither the router ID nor system interface are defined, the router ID from the base router context is inherited.

The **no** form of the command removes the router ID definition from the given VPRN context.

Default **no router-id**

Parameters *ip-address* — The IP address must be given in dotted decimal notation.

snmp-community

Syntax **snmp-community** *community-name* [**version** *SNMP-version*]
no snmp-community [*community-name*]

Context config>service>vprn

Description This command sets the SNMP community name to be used with the associated VPRN instance.

If an SNMP community name is not specified, then SNMP access is not allowed.

The **no** form of the command removes the SNMP community name from the given VPRN context.

Default None — The SNMP community must be explicitly specified.

Parameters *community-name* — Specify one or more SNMP community names.

version *SNMP-version* **—** Specify the SNMP version.

Values **v1**, **v2c**, **both**

source-address

Syntax **source-address**
no source-address

Context config>service>vprn

Description This command enters the context to specify the source address and application that should be used in all unsolicited packets.

application

Syntax **application** *app* [*ip-int-name*|*ip-address*]
no application *app*

Context config>service>vprn>source-address

Description This command specifies the source address and application.

Parameters *app* — Specify the application name.

Values telnet, ssh, traceroute, ping

ip-int-name|*ip-address* — Specifies the name of the IP interface or IP address. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

static-route

Syntax [**no**] **static-route**{*ip-prefix*/*mask* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **next-hop** [*ip-int-name* | *ip-address*] [**bfd-enable**]
[**no**] **static-route** {*ip-prefix*/*mask* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **indirect** *ip-address*
[**no**] **static-route** {*ip-prefix*/*mask* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **black-hole**

Context config>service>vprn

Description This command creates static route entries for network and access routes. To reduce entering configurations manually, address aggregation should be applied where possible. When configuring a static route, either next-hop, indirect, or black-hole must be configured.

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.

Default None — A static route must be explicitely defined.

Parameters *ip-prefix* — The destination address of the static route in dotted decimal notation.

mask — The mask associated with the network address, expressed as a mask length or in dotted decimal notation. For example, enter /16 for a 16-bit mask. The mask can also be entered in dotted decmial notation (255.255.0.0).

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.

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 metrices of an existing static route, unless specified, the preference will not change.

If there are multiple static routes with the same preference but unequal metrices, 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.

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.

type

Syntax **type** *[hub |subscriber-split-horizon]*
no type

Context config>service>vprn>

Description This command designates the type of VPRN instance being configured for hub and spoke topologies. Use the **no** form to reset to the default of a fully meshed VPRN.

Default **no type**

Parameters *hub* — Specifies a hub VPRN which allows all traffic from the hub SAPs to be routed to the destination directly, while all traffic from spoke VPRNs or network interfaces can only be routed to a hub SAP.

subscriber-split-horizon — Controls the flow of traffic for wholesale subscriber applications.

vrf-export

Syntax **vrf-export** *policy* [*policy...*]
no vrf-export

Context config>service>vprn

Description This command specifies the export policies to control routes exported from the local VPN routing/forwarding (VRF) to other VRFs on the same or remote PE routers (via MP-BGP).

The **no** form of the command removes all route policy names from the export list.

Default None **—** No routes are exported from the VRF by default.

Parameters *policy* — The route policy statement name.

vrf-import

Syntax **vrf-import** *policy* [*policy...*]
no vrf-import

Context config>service>vprn

Description This command sets the import policies to control routes imported to the local VPN routing/forwarding (VRF) from other VRFs on the same or remote PE routers (via MP-BGP). BGP-VPN routes imported with a vrf-import policy will use the BGP preference value of 170 when imported from remote PE routers, or retain the protocol preference value of the exported route when imported from other VRFs on the same router, unless the preference is changed by the policy.

The **no** form of the command removes all route policy names from the import list

Default None **—** No routes are accepted into the VRF by default.

Parameters *policy* — The route policy statement name.

vrf-target

vrf-target {*ext-comm* **|** [**import** *ext-comm*] [**export** *ext-comm*]}
no vrf-target

Context config>service>vprn

Description This command facilitates a simplified method to configure the route target to be added to advertised routes or compared against received routes from other VRFs on the same or remote PE routers (via MP-BGP).

BGP-VPN routes imported with a vrf-target statement will use the BGP preference value of 170 when imported from remote PE routers, or retain the protocol preference value of the exported route when imported from other VRFs in the same router.

Specified **vrf-import** or **vrf-export** policies override the **vrf-target** policy.

The no form of the command removes the vrf-target

Default **no vrf-target**

Parameters *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* |*as-number*:*ext-comm-val*}
ip-address: a.b.c.d
comm-val: 0 — 65535
as-number: 1 — 65535
ext-comm-val: 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.

Redundant Interface Commands

redundant-interface

Syntax [**no**] **redundant-interface** *ip-int-name*

Context config>service>vprn

Description This command configures a redundant interface.

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

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*} [**remote-ip** *ip-address*]
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 *ip-address/mask* — Assigns an IP address/IP subnet format to the interface.

ip-address 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.

Assigns an IP address netmask to the interface.

remote-ip *ip-address* — Assigns a remote IP to the interface.

SDP Commands

spoke-sdp

Syntax [**no**] **spoke-sdp** *sdp-id*

Context config>service>vprn

Description This command binds a service to an existing Service Distribution Point (SDP).

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

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a VPRN service. If the **sdp** *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 7750 SR devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

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

Special Cases **VPRN —** Several SDPs can be bound to a VPRN service. Each SDP must be destined to a different 7750 SR router. If two *sdp-id* bindings terminate on the same 7750 SR, an error occurs and the second SDP binding is rejected.

Parameters *sdp-id* — The SDP identifier. Allowed values are integers in the range of 1 and 17407 for existing SDPs.

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

spoke-sdp

Syntax **spoke-sdp** *sdp-id:vc-id*
no spoke-sdp *sdp-id:vc-id*

Context config>service>vprn>interface

Description This command binds a service to an existing Service Distribution Point (SDP).

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

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end 7750 SR devices can participate in the service.

The no form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Default No SDP 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 7750 SR router. If two sdp-id bindings terminate on the same 7750 SR, an error occurs and the second SDP is binding is rejected.

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

egress

Syntax **egress**

Context cconfig>service>vprn>if>spoke-sdp
config>service>vprn>red-if>spoke-sdp

Description This command configures an SDP context.

ingress

Syntax **ingress**

Context config>service>vprn>if>spoke-sdp
config>service>vprn>red-if>spoke-sdp

Description This command configures the SDP context.

vc-label

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

Context config>service>vprn>if>spoke-sdp>egress
config>service>vprn>red-if>spoke-sdp>egress

Description This command configures the egress VC label.

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

Values 16 — 1048575

vc-label

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

Context config>service>vprn>>if>spoke-sdp>ingress
config>service>vprn>red-if>spoke-sdp>ingress

Description This command configures the ingress VC label.

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

Values 2048 — 18431

filter

Syntax **filter** {**ip** *ip-filter-i*d}
no filter

Context 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

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

Interface Commands

interface

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

Context config>service>vprn

Description This command creates a logical IP routing interface for a Virtual Private Routed Network (VPRN). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.

The **interface** command, under the context of services, is used to create and maintain IP routing interfaces within VPRN service IDs. The **interface** command can be executed in the context of an VPRN service ID. The IP interface created is associated with the service core network routing instance and default routing table. The typical use for IP interfaces created in this manner is for subscriber internet access.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config router interface** and **config service vprn interface**. Interface names must not be in the dotted decimal notation of an IP address. For example, the name "1.1.1.1" is not allowed, but "int-1.1.1.1" is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

The available IP address space for local subnets and routes is controlled with the **config router service-prefix** command. The **service-prefix** command administers the allowed subnets that can be defined on service IP interfaces. It also controls the prefixes that may be learned or statically defined with the service IP interface as the egress interface. This allows segmenting the IP address space into **config router** and **config service** domains.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All VPRN IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

The **no** form of this command removes IP the interface and all the associated configuration. The interface must be administratively shutdown before issuing the **no interface** command.

For VPRN services, the IP interface must be shutdown before the SAP on that interface may be removed. VPRN services do not have the **shutdown** command in the SAP CLI context. VPRN service SAPs rely on the interface status to enable and disable them.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and config service vprn interface commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, $, spaces, 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.

active-cpm-protocols

Syntax [**no**] **active-cpm-protocols**

Context config>service>vprn>interface

Description This command enables CPM protocols on this interface.

address

Syntax **address** *ip-address/mask* | *ip-address netmask*} [**broadcast** [**all-ones** | **host-ones**]
no address

Context config>service>vprn>interface

Description Assigns an IP address, IP subnet, and broadcast address format to a VPRN IP router interface. Only one IP address can be associated with an IP interface.

An IP address must be assigned to each VPRN IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the 7750 SR.

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

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

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

Use the **no** form of this command to remove the IP address assignment from the IP interface. When the **no address** command is entered, the interface becomes operationally down.

Address	Admin state	Oper state
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an adminstratively 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 /30 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 indictates 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.

allow-directed-broadcasts

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

Context config>service>vprn>interface

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

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

bfd

Syntax **bfd** *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier*]
no bfd

Context config>service>vprn>interface

Description This command specifies the BFD parameters for the associated IP interface. If no parameters are defined the default value are used.

The multiplier specifies the number of consecutive BFD messages that must be missed from the peer before the BFD session state is changed to down and the upper level protocols (OSPF, IS-IS or PIM) is notified of the fault.

The **no** form of the command removes BFD from the associated IGP protocol adjacency.

Default **no bfd**

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

Values 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

cflowd

Syntax **cflowd** {**acl** | **interface**}
no cflowd

Context config>service>vprn>interface

Description This command enables **cflowd** to collect traffic flow samples through a router for analysis.

cflowd is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement. When **cflowd** is enabled at the interface level, all packets forwarded by the interface are subjected to analysis according to the **cflowd** configuration.

Default **no cflowd**

Parameters *ACL — cflowd* configuration associated with a filter.

interface — cflowd configuration associated with an IP interface.

ip-mtu

Syntax **ip-mtu** *octets*
no ip-mtu

Context config>service>vprn>interface

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

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

Default **no ip-mtu**

local-proxy-arp

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

Context config>service>vprn>interface
config>service>vprn>sub-if>grp-if

Description This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus will become the forwarding point for all traffic between hosts in that subnet.

When local-proxy-arp is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

Default **no local-proxy-arp**

loopback

Syntax [**no**] **loopback**

Context config>service>vprn>interface
config>service>vprn>interface

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>interface
config>service>vprn>if>vrrp
config>service>vprn>sub-if>grp-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.

proxy-arp-policy

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

Context config>service>vprn>interface
config>service>vprn>sub-if>grp-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 *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.

redundant-interface

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

Context config>service>vprn
config>service>vprn>sub-if>grp-if

Description This command configures a redundant interface used for dual homing.

Parameters *red-ip-int-name* — Specifies the redundant IP interface name.

remote-proxy-arp

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

Context config>service>vprn>interface
config>service>vprn>sub-if>grp-if

Description This command enables remote proxy ARP on the interface.

Remote proxy ARP is similar to proxy ARP. It allows the router to answer an ARP request on an interface for a subnet that is not provisioned on that interface. This allows the router to forward to the other subnet on behalf of the requester. To distinguish remote proxy ARP from local proxy ARP, local proxy ARP performs a similar function but only when the requested IP is on the receiving interface.

Default **no remote-proxy-arp**

secondary

Syntax **secondary** {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**] [**igp-inhibit**]
no secondary {*ip-address/mask* | *ip-address netmask*}

Context config>service>vprn>interface

Description This command assigns an secondary IP address/IP subnet/broadcast address format to the interface.

Default **none**

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 /30 subnets).

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical ‘AND’ function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. 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 indictates 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.

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 [**no**] **static-arp** *ip-address ieee-mac-address*

Context config>service>vprn>interface

Description This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP will appear in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.

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

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

Default No static ARPs defined.

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.

tos-marking-state

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

Context config>service>vprn>interface
config>service>vprn>sub-if>grp-if

Description This command is used to alter the default trusted state to a non-trusted state. When unset or reverted to the trusted default, the ToS field will not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set, in which case the egress network interface treats all VPRN and network IP interface as untrusted.

When the ingress interface is set to untrusted, all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface. The egress network remarking rules also apply to the ToS field of IP packets routed using IGP shortcuts (tunneled to a remote next-hop). However, the tunnel QoS markings are always derived from the egress network QoS definitions.

Egress marking and remarking is based on the internal forwarding class and profile state of the packet once it reaches the egress interface. The forwarding class is derived from ingress classification functions. The profile of a packet is either derived from ingress classification or ingress policing.

The default marking state for network IP interfaces is trusted. This is equivalent to declaring no tos-marking-state on the network IP interface. When undefined or set to tos-marking-state trusted, the trusted state of the interface will not be displayed when using show config or show info unless the detail parameter is given. The **save config** command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified.

The **no** tos-marking-state command is used to restore the trusted state to a network IP interface. This is equivalent to executing the tos-marking-state trusted command.

Default **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>interface

Description This command configures the interface as an unnumbered interface.

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.

DHCP Commands

dhcp

Syntax **dhcp**

Context config>service>vprn>interface
config>service>vprn>subscriber-interface
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure DHCP parameters.

action

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

Context config>service>vprn>if>dhcp>option
config>service>vprn>sub-if>grp-if>dhcp>option

Description This command configures the processing required when the SR-Series receives a DHCP request that already has a Relay Agent Information Option (Option 82) field in the packet.

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

Default Per RFC 3046, *DHCP Relay Agent Information Option* , section 2.1.1, *Reforwarded DHCP requests*, the default is to keep the existing information intact. The exception to this is if the giaddr of the received packet is the same as the ingress address on the router. In that case the packet is dropped and an error is logged.

Parameters **replace —** In the upstream direction (from the user), the existing Option 82 field is replaced with the Option 82 field from the router. In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop — The packet is dropped, and an error is logged.

keep — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is sent on towards the client.

The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.

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

circuit-id

Syntax **circuit-id** [**ascii-tuple** | **ifindex** | **sap-id** | **vlan-ascii-tuple**]
no circuit-id

Context config>service>vprn>if>dhcp>option
config>service>vprn>sub-if>grp-if>dhcp>option

Description When enabled, the router sends the interface index (If Index) in the **circuit-id** suboption of the DHCP packet. The If Index of a router interface can be displayed using the command `show>router>interface>detail`. This option specifies data that must be unique to the router that is relaying the circuit.

If disabled, the **circuit-id** suboption of the DHCP packet will be left empty.

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

Default **circuit-id**

Parameters **ascii-tuple —** Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by "|".

ifindex — Specifies that the interface index will be used. The If Index of a router interface can be displayed using the command `show>router>interface>detail`.

sap-id — Specifies that the SAP ID will be used.

vlan-ascii-tuple — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

gi-address

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

Context config>service>vprn>if>dhcp
config>service>vprn>sub-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp

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

Default **no gi-address**

Parameters *ip-address* — Specifies the host IP address to be used for DHCP relay packets.

src-ip-address — Specifies the source IP address to be used for DHCP relay packets.

lease-populate

Syntax **lease-populate** [*nbr-of-entries*]
no lease-populate

Context config>service>vprn>if>dhcp
config>service>vprn>sub-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp

Description This command enables dynamic host lease state management for VPLS SAPs and VPRN or IES IP interfaces. Lease state information is extracted from snooped or relayed DHCP ACK messages to populate lease state table entries for the SAP or IP interface.

The optional *number-of-entries* parameter defines the number lease state table entries allowed for this SAP or IP interface. If *number-of-entries* is not specified, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed. If lease state population is enabled and an entry cannot be retained in the table, the DHCP Relay or DHCP snoop function will prevent the far-end host from receiving the DHCP ACK message.

The retained lease state information representing dynamic hosts may be used to populate a SAP based anti-spoof filter table to provide dynamic anti-spoof filtering. Anti-spoof filtering is only available on VPLS SAPs, IES IP interfaces terminated on a SAP or VPRN IP interfaces terminated on a SAP.

The retained lease state information representing dynamic hosts may be used to populate the system's ARP cache based the arp-populate feature. ARP-populate functionality is only available for static and dynamic hosts associated with IES and VPRN SAP bound IP interfaces.

The retained lease state information representing dynamic hosts may be used to populate managed entries into a VPLS forwarding database. VPLS forwarding database population is an implicit feature that automatically places the dynamic host's MAC address into the VPLS FDB. When a dynamic host's MAC address is placed in the lease state table, it will automatically be populated into the VPLS forwarding database associated with the SAP on which the host is learned. The dynamic host MAC address will override any static MAC entries using the same MAC and prevent dynamic learning of the MAC on another interface. Existing static MAC entries with the same MAC address as the dynamic host are marked as inactive but not deleted. If all entries in the lease state table associated with the MAC address are removed, the static MAC may be populated. New static MAC definitions for the VPLS instance may be created while a dynamic host exists associated with the static MAC address.

Default **not enabled**

Parameters *nbr-of-entries* — Defines the number lease state table entries allowed for this interface. If this parameter is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCP ACK messages are discarded.

Values 1 — 8000

match-circuit-id

Syntax [**no**] **match-circuit-id**

Context config>service>vprn>sub-if>grp-if>dhcp

Description This command enables Option 82 circuit ID on relayed DHCP packet matching.

For Routed CO, the group interface DHCP relay process is stateful. When packets are relayed to the server the virtual router ID, transaction ID, SAP ID, and client hardware MAC address of the relayed packet are tracked. When we get a response back from the server the virtual router ID, transaction ID, and client HW MAC address must be matched to determine the SAP on which to send the packet out. In some cases, the virtual router ID, transaction ID, and client HW MAC address are not guaranteed to be unique.

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

Default **no match-circuit-id**

option

Syntax [**no**] **option**

Context config>service>vprn>if>dhcp
config>service>vprn>sub-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp

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

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

Default **no option**

remote-id

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

Context config>service>vprn>sub-if>grp-if>dhcp>option

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

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

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

Default **remote-id**

Parameters **mac —** This keyword specifies the MAC address of the remote end is encoded in the suboption.

string *string* **—** Specifies the remote-id.

vendor-specific-option

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

Context config>service>vprn>if>dhcp>option
config>service>vprn>sub-if>grp-if>dhcp>option

Description This command configures the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

client-mac-address

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

Context config>service>vprn>if>dhcp>option
config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

Description This command enables the sending of the MAC address in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

sap-id

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

Context config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

Description This command enables the sending of the SAP ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the SAP ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

service-id

Syntax [**no**] **service-id**

Context config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

Description This command enables the sending of the service ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

string

Syntax [**no**] **string** *text*

Context config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

Description This command specifies the string in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

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

Parameters *text* — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

system-id

Syntax [**no**] **system-id**

Context config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

Description This command specifies whether the system-id is encoded in the Alcatel-Lucent vendor specific sub-option of Option 82.

Default **None**

proxy-server

Syntax **proxy-server**

Context config>service>if>dhcp
config>service>vprn>sub-if>grp-if>dhcp

Description This command configures the DHCP proxy server.

emulated-server

Syntax **emulated-server** *ip-address*
no emulated-server

Context config>service>vprn>if>dhcp>proxy
config>service>vprn>sub-if>grp-if>dhcp>proxy-server

Description This command configures the IP address to be used as the DHCP server address in the context of this service. Typically, the configured address should be in the context of the subnet.

The **no** form of this command reverts to the default setting. The local proxy server will not become operational without a specified emulated server address.

Parameters *ip-address* — Specifies the emulated server address.

Default Note that for a retail interface, the default is the local interface.

lease-time

Syntax **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*] [**radius-override**]
no lease-time

Context config>service>vprn>if>dhcp>proxy
config>service>vprn>sub-if>grp-if>dhcp>proxy-server

Description This command defines the length of lease-time that will be provided to DHCP clients. By default the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.

The no form of this command disables the use of the lease-time command. The local-proxy-server will use the lease-time offered by either a RADIUS or DHCP server.

Default 7 days 0 hours 0 seconds

Parameters **radius-override —** Specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients.

days — Specifies the number of days that the given IP address is valid.

Values 0 — 3650

hours — Specifies the number of hours that the given IP address is valid.

Values 0 — 23

minutes — Specifies the number of minutes that the given IP address is valid.

Values 0 — 59

seconds — Specifies the number of seconds that the given IP address is valid.

Values 0 — 59

server

Syntax **server** *server1* [*server2*...(up to 8 max)]

Context config>service>vprn>if>dhcp
config>service>vprn>sub-if>grp-if>dhcp

Description This command specifies a list of servers where requests will be forwarded. The list of servers can entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list.

There can be a maximum of 8 DHCP servers configured.

The flood command is applicable only in the VPLS case. There is a scenario with VPLS where the VPLS node only wants to add Option 82 information to the DHCP Request to provider per-subscriber information, but it does not do full DHCP Relay. In this case, the server is set to "flood". This means the DHCP Request is still a broadcast and is sent through the VPLS domain. A node running at L3 further upstream then can perform the full L3 DHCP Relay function.

Default **no server**

Parameters *server* — Specify the DHCP server IP address.

trusted

Syntax [**no**] **trusted**

Context config>service>vprn>if>dhcp

config>service>vprn>sub-if>grp-if>dhcp

Description According to RFC 3046, *DHCP Relay Agent Information Option*, a DHCP request where the giaddr is 0.0.0.0 and which contains a Option 82 field in the packet, should be discarded, unless it arrives on a "trusted" circuit. If trusted mode is enabled on an IP interface, the relay agent (the SR-Series) will modify the request's giaddr to be equal to the ingress interface and forward the request.

Note that this behavior only applies when the action in the Relay Agent Information Option is "keep". In the case where the Option 82 field is being replaced by the relay agent (action = "replace"), the original Option 82 information is lost anyway, and there is thus no reason for enabling the trusted option.

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

Default **not enabled**

host-connectivity-verify

Syntax **host-connectivity-verify** [**interval** *interval*] [**action** {**remove**|**alarm**}]

Context config>service>vprn>if>sap
config>service>vprn>sub-if>grp-if
config>service>vprn>sub-if>grp-if>dhcp

Description This command enables 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 **interval *interval* —** The interval, expressed in minutes, which specifies the time interval which all known sources should be verified. The actual rate is then dependent on number of known hosts and interval.

Values 1— 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, 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.

Subscriber Interface Commands

subscriber-interface

Syntax **subscriber-interface** *ip-int-name* [**fwd-service** *service-id* **fwd-subscriber-interface** *ip-int-name*]
no subscriber-interface *ip-int-name*

Context config>service>vprn

Description This command allows the operator to create special subscriber-based interfaces. It is used to contain multiple group interfaces. Multiple subnets associated with the subscriber interface can be applied to any of the contained group interfaces in any combination. The subscriber interface allows subnet sharing between group interfaces.

Use the **no** form of the command to remove the subscriber interface.

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

fwd-service *service-id* **—** Specifies the forwarding service ID for a subscriber interface in a retailer context.

fwd-subscriber-interface *ip-int-name* **—** Specifies the forwarding subscriber interface for a subscriber interface in a retailer context.

address

Syntax [**no**] **address** {*ip-address/mask* | *ip-address netmask*} [**gw-ip-address** *ip-address*]

Context config>service>vprn>subscriber-interface

Description This command configures the local subscriber subnets available on a subscriber IP interface. The configured ip-address and mask define the address space associated with the subscriber subnet. Up to 16 IP subnets can be created on a single subscriber IP interface. Each subnet supports a locally owned IP host address within the subnet that is not expected to appear on other routers that may be servicing the same subscriber subnet. For redundancy purposes, the keyword **gw-address** defines a separate IP address within the subnet for Subscriber Routed Redundancy Protocol (SRRP) routing. This IP address must be the same on the local and remote routers participating in a common SRRP instance.

In SRRP, a single SRRP instance is tied to a group IP interface. The group IP interface is contained directly within a subscriber IP interface context and thus directly associated with the subscriber subnets on the subscriber IP interface. The SRRP instance is also indirectly associated with any subscriber subnets tied to the subscriber interface through wholesale/retail VPRN configurations. With the directly-associated and the indirectly-associated subscriber interface subnets, a single SRRP instance can manage hundreds of SRRP gateway IP addresses. This automatic subnet association to the SRRP instance is different from VRRP where the redundant IP address is defined within the VRRP context.

Defining an SRRP gateway IP address on a subscriber subnet is not optional when the subnet is associated with a group IP interface with SRRP enabled. Enabling SRRP (**no shutdown**) will fail if

one or more subscriber subnets do not have an SRRP gateway IP address defined. Creating a new subscriber subnet without an SRRP gateway IP address defined will fail when the subscriber subnet is associated with a group IP interface with an active SRRP instance. Once SRRP is enabled on a group interface, the SRRP instance will manage the ARP response and routing behavior for all subscriber hosts reachable through the group IP interface.

The no form of the command removes the address from a subscriber subnet. The address command for the specific subscriber subnet must be executed without the gw-address parameter. To succeed, all SRRP instances associated with the subscriber subnet must removed or shutdown.

Parameters *ip-address/mask* | *ip-address netmask* — Specifies the address space associated with the subscriber subnet

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

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

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

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

group-interface

Syntax [**no**] **group-interface** *ip-int-name*

Context config>service>vprn>sub-if
config>service>vprn>subscriber-interface

Description This command enables the context to configure a group interface. A group interface is an interface that may contain one or more SAPs. This interface is used in triple-play services where multiple SAPs are part of the same subnet.

Default **none**

Parameters *ip-int-name* — Configures the interface group name. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Interface ICMP Commands

icmp

Syntax **icmp**

Context config>service>vprn>interface
config>service>vprn>sub-if>grp-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>sub-if>grp-if>icmp

Description This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance will reply to mask requests.

The **no** form of this command disables replies to ICMP mask requests on the router interface.

Default **mask-reply** - Reply to ICMP mask requests.

redirects

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

Context config>service>vprn>if>icmp
config>service>vprn>sub-if>grp-if>icmp

Description This commad configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.

When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.

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

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

The **no** form of this command disables the generation of icmp redirects on the router interface.

Default **redirects 100 10** - 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>sub-if>grp-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

unreachables

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

Context config>service>vprn>if>icmp
config>service>vprn>sub-if>grp-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 SAP Commands

sap

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

Context config>service>vprn>interface
config>service>vprn>sub-if>grp-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 7750. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface** *port-type port-id* **mode access** command. Channelized TDM ports are always access ports.

If a port is 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.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP will also be deleted.

Default No SAPs are defined.

Special Cases **VPRN —** A VPRN SAP must be defined on an Ethernet interface.

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

The *sap-id* can be configured in one of the following formats:

Type	Syntax	Example
port-id	*slot*/*mda*/*port*[.*channel*]	2/2/11 1/2/3.1
null	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id* | *aps-id*]	*port-id*: 1/2/3 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*: lag-100 *aps-id*: aps-1
dot1q	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id* | *aps-id*]:qtag1	*port-id*:qtag1: 1/2/3:100 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*:qtag1:lag-100:102 *aps-id*:qtag1: aps-1:103

Type	Syntax	Example
qinq	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	*port-id*:qtag1.qtag2: 1/2/3:100.10 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*:qtag1.qtag2:lag-100:
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:vpi/vci |vpi |vpi1.vpi2]	port-id: 9/1/1 aps-id: aps-1 *bundle-id*: bundle-ima-1/1.1 bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 vpi/vci: 16/26 vpi: 16 vpi1.vpi2: 16.200
frame-relay	[*port-id* | *aps-id*]:*dlci*	*port-id*: 1/1/1:100 *aps-id*: aps-1 *dlci*: 16
cisco-hdlc	*slot/mda/port.channel*	*port-id*: 1/2/3.1

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200

qtag1	0 — 4094
qtag2	*, 0 — 4094

vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

bundle-id — Specifies the multilink bundle to be associated with this IP interface. The **bundle** keyword must be entered at the beginning of the parameter.
The command syntax must be configured as follows:

bundle-id: **bundle**-*type-slot-id*/*mda-slot*.*bundle-num*
bundle-id value range: 1 — 128

For example:

```
*A:ALA-12>config#  port bundle-ppp-5/1.1
*A:ALA-12>config>port# multilink-bundle
```

bgprp-id — Specifies the bundle protection group ID to be associated with this IP interface. The **bpgrp** keyword must be entered at the beginning of the parameter.

The command syntax must be configured as follows:

bpgrp-id: **bpgrp**-*type-bpgrp-num*
type: ima
bpgrp-num value range: 1 — 1280

For example:

```
*A:ALA-12>config# port bpgrp-ima-1
*A:ALA-12>config>service>vpls$ sap bpgrp-ima-1
```

qtag1, *qtag2* — Specifies the encapsulation value used to identify the SAP on the port or sub-port. If this parameter is not specificially defined, the default value is 0.

Values

qtag1:	0 — 4094
qtag2 :	* | 0 — 4094

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 TDM	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 TDM	BCP-Dot1q	0 — 4094	The SAP is identified by the 802.1Q tag on the channel.
SONET/SDH TDM	Frame Relay	16 — 991	The SAP is identified by the data link connection identifier (DLCI).
SONET/SDH ATM	ATM	vpi (NNI) 0 — 4095 vpi (UNI) 0 — 255 vci 1, 2, 5 — 65535 -	The SAP is identified by port or by PVPC or PVCC identifier (vpi, vpi/vci, or vpi range)

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.

tod-suite

Syntax **tod-suite** *tod-suite-name*
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>cron** context.

Default **no tod-suite**

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

Interface Billing Commands

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
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 *policy-id* does not exist, an error message is generated.

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

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

Default Default accounting policy.

Parameters *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 to 99

collect-stats

Syntax [**no**] **collect-stats**

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

Description This command enables accounting and statistical data collection for either an interface SAP or interface SAP spoke SDP, or network port. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default **collect-stats**

Interface SAP ATM Commands

atm

Syntax **atm**

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

Description This command enables access to the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as:

- Configuring ATM port or ATM port-related functionality on MDAs supporting ATM functionality
- Configuring ATM-related configuration for ATM-based SAPs that exist on MDAs supporting ATM functionality.

If ATM functionality is not supported for a given context, the command returns an error.

egress

Syntax **egress**

Context config>service>vprn>if>sap>atm
config>service>vprn>sub-if>grp-if>sap>atm

Description This command configures egress ATM attributes for the SAP.

encapsulation

Syntax **encapsulation** *atm-encap-type*

Context config>service>vprn>if>sap>atm
config>service>vprn>sub-if>grp-if>sap>atm

Description This command configures RFC 2684, *Multiprotocol Encapsulation over ATM AAL5*, encapsulation for an ATM PVCC delimited SAP.

This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition references RFC 2684 and to the ATM Forum LAN Emulation specification.

Ingress traffic that does not match the configured encapsulation will be dropped.

Default The encapsulation is driven by the services for which the SAP is configured.
For VPRN service SAPs, the default is **aal5snap-routed**.

Parameters *atm-encap-type* — Specify the encapsulation type.

Values **aal5snap-routed** — Routed encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684.

aal5mux-ip — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684.
aal5snap-bridged — Bridged encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684.
aal5mux-bridged-eth-nofcs — Bridged IP encapsulation for VC multiplexed circuit as defined in RFC 2684.

ingress

Syntax **ingress**

Context config>service>vprn>if>sap>atm
config>service>vprn>sub-if>grp-if>sap>atm

Description This command configures ingress ATM attributes for the SAP.

traffic-desc

Syntax **traffic-desc** *traffic-desc-profile-id*
no traffic-desc

Context config>service>vprn>if>sap>atm>egress
config>service>vprn>if>sap>atm>ingress
config>service>vprn>sub-if>grp-if>sap>atm>egress
config>service>vprn>sub-if>grp-if>sap>atm>ingress

Description This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP).

When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction.

When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.

The **no** form of the command reverts the traffic descriptor to the default traffic descriptor profile.

Default The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.

Parameters *traffic-desc-profile-id* — Specify a defined traffic descriptor profile (see the QoS atm-td-profile command).

oam

Syntax **oam**

Context config>service>vprn>interface >sap>atm
config>service>vprn>sub-if>grp-if>sap>atm

Description This command enables the context to configure OAM functionality for a PVCC delimiting a SAP.

The ATM-capable MDAs support F5 end-to-end OAM functionality (AIS, RDI, Loopback):

- ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95
- GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996
- GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

alarm-cells

Syntax [**no**] **alarm-cells**

Context config>service>vprn>if>sap>atm>oam
config>service>vprn>sub-if>grp-if>sap>atm>oam

Description This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC termination to monitor and report the status of their connection by propagating fault information through the network and by driving PVCC's operational status.

When alarm-cells functionality is enabled, a PVCC's operational status is affected when a PVCC goes into an AIS or RDI state because of an AIS/RDI processing (assuming nothing else affects PVCC's operational status, for example, if the PVCC goes DOWN, or enters a fault state and comes back UP, or exits that fault state). RDI cells are generated when PVCC is operationally DOWN. No OAM-specific SNMP trap is raised whenever an endpoint enters/exits an AIS or RDI state, however, if as result of an OAM state change, the PVCC changes operational status, then a trap is expected from an entity the PVCC is associated with (for example a SAP).

The **no** command disables alarm-cells functionality for a PVCC. When alarm-cells functionality is disabled, a PVCC's operational status is no longer affected by a PVCC's OAM state changes due to AIS/RDI processing (note that when alarm-cells is disabled, a PVCC will change operational status to UP due to alarm-cell processing) and RDI cells are not generated as result of the PVCC going into AIS or RDI state. The PVCC's OAM status, however, will record OAM faults as described above.

Default Enabled for PVCCs delimiting VPRN SAPs

periodic-loopback

Syntax [**no**] **periodic-loopback**

Context config>service>vprn>if >sap>atm>oam
config>service>vprn>sub-if>grp-if>sap>atm

Description This command enables periodic OAM loopbacks on this SAP. This command is only configurable on VPRN SAPs. When enabled, an ATM OAM loopback cell is transmitted every period as configured in the `config>system>atm>oam>loopback-period` *`period`* context.

If a response is not received and consecutive retry-down retries also result in failure, the endpoint will transition to an alarm indication signal/loss of clock state. Then, an ATM OAM loopback cell will be transmitted every period as configured in the `loopback-period` *`period`*. If a response is received for the periodic loopback and consecutive retry-up retries also each receive a response, the endpoint will transition back to the up state.

The **no** form of the command sets the value back to the default.

Default **no periodic-loopback**

Interface Anti-Spoofing Commands

anti-spoof

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

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

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

The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (**ip**, **mac**, **ip-mac**) is 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 Filter type default types:

- 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 **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.

arp-populate

Syntax [**no**] **arp-populate**

Context config>service>vprn>if
config>service>vprn>sub-if>subscriber-interface
config>service>vprn>sub-if>grp-if

Description This command enables populating static and dynamic hosts into the system ARP cache. When enabled, the host's IP address and MAC address are placed in the system ARP cache as a managed

entry. Static hosts must be defined on the interface using the **host** command. Dynamic hosts are enabled on the system through enabling lease-populate in the IP interface DHCP context. In the event that both a static host and a dynamic host share the same IP and MAC address, the system's ARP cache retains the host information until both the static and dynamic information are removed. Both static and dynamic hosts override static ARP entries. Static ARP entries are marked as inactive when they conflict with static or dynamic hosts and will be repopulated once all static and dynamic host information for the IP address are removed. Since static ARP entries are not possible when static subscriber hosts are defined or when DHCP lease state table population is enabled, conflict between static ARP entries and the arp-populate function is not an issue.

The **arp-populate** command 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.

arp-populate can only be enabled on VPRN interfaces supporting Ethernet encapsulation.

Use the **no** form of the command to disable ARP cache population functions for static and dynamic hosts on the interface. All static and dynamic host information in the systems ARP cache will be removed. Any existing static ARP entries previously inactive due to static or dynamic hosts will be populated in the system ARP cache.

When **arp-populate** is enabled, the system will not send out ARP Requests for hosts that are not in the ARP cache. Only statically configured and DHCP learned hosts are reachable through an IP interface with arp-populate enabled.

Default **not enabled**

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>vprn>interface
config>service>vprn>sub-if>grp-if

Description This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

The default value for **arp-timeout** is 14400 seconds (4 hours).

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** *name*
no authentication-policy

Context config>service>vprn>if
config>service>vprn>sub-if>grp-if

Description This command assigns an authentication policy to the interface.

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

Default **no authentication-policy**

Parameters *name* — Specifies the authentication policy name. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

host

Syntax [**no**] **host** {[**ip** *ip-address* [**mac** *ieee-address*]} [**subscriber** *sub-ident-string*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*]
no host {[**ip** *ip-address*] [**mac** *ieee-address*]}

Context config>service>vprn>if>sap

Description This command creates a static host for the SAP. Applications within the system that make use of static host entries include anti-spoof, and source MAC population into the VPLS forwarding database.

Multiple static hosts can be defined on the SAP. Each host is identified by a source IP address, a source MAC address, or both a source IP and source MAC address. When anti-spoof in enabled on the SAP, the host information will be populated into the SAP's anti-spoof table, allowing ingress packets matching the entry access to the SAP. When the MAC address exists in the host definition, the MAC address is populated into the VPLS forwarding database and associates it with the SAP. The static host definition overrides any static MAC entries using the same MAC and prevents dynamic learning of the MAC on another interface.

Defining a static host identical to an existing static host has no effect and will not generate a log or error message.

Every static host definition must have at least one address defined, IP or MAC.

Static hosts may exist on the SAP even with anti-spoof and **arp-populate** (VPRN) features disabled. When enabled, each feature has different requirements for static hosts.

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

arp-populate — When enabled, this feature uses static and dynamic host information to populate entries into the system's ARP cache. This is only available on the VPRN service SAPs. Both a

MAC address and IP address are required to populate an ARP entry in the system. If definition of a static host is attempted without both a MAC and IP address specified when arp -populate is enabled, the static host definition will fail.

fdb-populate — This is an implicit feature that uses the static host definition as a static MAC in the VPLS forwarding database. It cannot be enabled or disabled and has no effect on the ability to create static hosts without a MAC address specified. When a MAC address is specified for a static host, it will automatically be populated into the VPLS forwarding database associated with the SAP on which the host is created. The static host MAC address will override any static MAC entries using the same MAC and prevent dynamic learning of the MAC on another interface. Existing static MAC entries with the same MAC address as a static host are marked as inactive but not deleted. If all static hosts are removed from the SAP, the static MAC may be populated. New static MAC definitions for the VPLS instance may be created while a static host exists associated with the static MAC address.

The **no** form of the command removes a static entry from the system. The specified **ip** *address* and **mac** *address* must match the host's exact IP and MAC addresses as defined when it was created. When a static host is removed from the SAP, the affect of its removal on the anti-spoof filter, ARP cache or the VPLS forwarding database is also evaluated.

Default There are no default static entries.

Parameters **ip** ***ip-address*** **—** Specify this optional parameter when defining a static host. The IP address must be specified for **anti-spoof** *ip* and **anti-spoof** *ip-mac* commands. Only one static host can be configured on the SAP with a given IP address.
The following rules apply to configure static hosts using an IP address:

- Only one static host can be defined using a specific IP address.
- Defining a static host with the same IP address as a previous static host overwrites the previous static host.
- If a static host has an IP address assigned, the MAC address for the host is optional (depending on the features enabled on the SAP).

mac mac-address **—** Specify this optional parameter when defining a static host. The MAC address must be specified for **anti-spoof** *mac* and **anti-spoof** *ip-mac*. Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address. The following rules apply to configuring static hosts using a MAC address:

- Multiple static hosts may share the same MAC address.
- Executing the host command with the same MAC address but a different IP address as an existing static host will create a new static host.
- If a static host has a MAC address assigned, the IP address for the host is optional (depending on the features enabled on the SAP).

Values 8k static and dynamic hosts per 10G forwarding complex. 64k per system.

subscriber *sub-ident-string* — Specify this optional parameter to specify an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

- For VPRN SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's sub-ident-string is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the

destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPRN destinations.

If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.

If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

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

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

host

Syntax **host** {[**ip** *ip-address*] [**mac** *mac-address*]}[**subscriber** *sub-ident-string*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**ancp-string** *ancp-string*]
no host {[**ip** *ip-address*] [**mac** *ieee-address*]}
no host all

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

Description This command creates a static subscriber host for the SAP. Static subscriber hosts may be used by the system for various purposes. Applications within the system that make use of static host entries include anti-spoof, ARP reply agent and source MAC population into the VPRN forwarding database.

Multiple static hosts may be defined on the SAP. Each host is identified by either a source IP address, a source MAC address or both a source IP and source MAC address. Every static host definition must have at least one address defined, IP or MAC.

Static hosts can exist on the SAP even with anti-spoof and ARP reply agent features disabled. When enabled, each feature has different requirements for static hosts.

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

Default **none**

Parameters **ip** *ip-address* — Specify this optional parameter when defining a static host. The IP address must be specified for **anti-spoof ip** and **anti-spoof ip-mac** and **arp-reply-agent**. Only one static host may be configured on the SAP with a given IP address.

mac *mac-address* — Specify this optional parameter when defining a static host. The MAC address must be specified for **anti-spoof mac**, and **anti-spoof ip-mac** and arp-reply-agent. Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address.

Every static host definition must have at least one address defined, IP or MAC.

subscriber *sub-ident-string* — Specify this optional parameter to specify an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

- For VPRN SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's sub-ident-string is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPRN destinations.

  If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

  If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.

  If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

  ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

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

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

all — Used with the **no** form of the command, the all keyword removes all **host** parameters specified in the **host** command.

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

Interface SAP Filter and QoS Policy Commands

egress

Syntax **egress**

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

Description This command enables the context to configure egress SAP Quality of Service (QoS) policies and filter policies.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing. If no egress filter is defined, no filtering is performed.

ingress

Syntax **ingress**

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

Description This command enables the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

agg-rate-limit

Syntax **agg-rate-limit** *agg-rate*
no agg-rate-limit

Context config>service>vprn>if>sap>egress
config>service>vprn>sub-if>grp-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 is bound to in order for the defined agg-rate-limit to take effect. The egress port scheduler enforces the aggregate queue rate as it distributes its bandwidth at the various port priority levels. The port

scheduler stops offering bandwidth to member queues once it has detected that the aggregate rate limit has been reached.

If a port scheduler is not defined on the egress port, the queues are allowed to operate based on their own bandwidth parameters.

The **no** form of the command removes the aggregate rate limit from the SAP or multi-service site.

Parameters *agg-rate* — defines the rate, in kilobits-per-second, that the maximum aggregate rate that the queues on the SAP or multi-service site can operate.

Values 1 — 40000000, max

filter

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

Context config>service>vprn>if>sap>egress
config>service>vprn>if>sap>ingress
config>service>vprn>sub-if>grp-if>sap>egress

Description This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface.

Filter policies control the forwarding and dropping of packets based on IP matching criteria.

The **filter** command is used to associate a filter policy with a specified *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

match-qinq-dot1p

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

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

Description This command specifies which Dot1Q tag position Dot1P bits in a QinQ encapsulated packet should be used to evaluate Dot1P QoS classification.

The **match-qinq-dot1p** command allows the top or bottom PBits to be used when evaluating the applied sap-ingress QoS policy's Dot1P entries. The **top** and **bottom** keywords specify which position should be evaluated for QinQ encapsulated packets.

The **no** form of the command restores the default dot1p evaluation behavior for the SAP.

By default, the bottom most service delineating Dot1Q tags Dot1P bits are used. Table 1 defines the default behavior for Dot1P evaluation when the **match-qinq-dot1p** command is not executed.

Table 1: 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 2 defines the dot1p evaluation behavior when the top parameter is specified.

Table 2: 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

Table 2: Top Position QinQ and TopQ SAP Dot1P Evaluation (Continued)

Port / SAP Type	Existing Packet Tags	PBits Used for Match
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 3 defines the dot1p evaluation behavior when the bottom parameter is specified.

Table 3: Bottom Position QinQ and TopQ SAP Dot1P Evaluation

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	BottomQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	BottomQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

Table 4: 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

Table 4: Default Dot1P Explicit Marking Actions (Continued)

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
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 5: QinQ Mark Top Only Explicit Marking Actions

Egress SAP Type	Ingress Packet Preserved Dot1P State	Marked (or Remarked) PBits
Null	No preserved Dot1P bits	None
Null	Preserved Dot1P bits	Preserved tag PBits remarked using dot1p-value
Dot1Q	No preserved Dot1P bits	New PBits marked using dot1p-value
Dot1Q	Preserved Dot1P bits	Preserved tag PBits remarked using dot1p-value
TopQ	No preserved Dot1P bits	TopQ PBits marked using dot1p-value
TopQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits marked using dot1p-value, BottomQ PBits preserved
QinQ	No preserved Dot1P bits	TopQ PBits PBits marked using dot1p-value, BottomQ PBits marked with zero
QinQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits PBits marked using dot1p-value, BottomQ PBits marked using preserved value

The QinQ and TopQ SAP PBit marking follows the default behavior devined in Table 4 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.

qinq-mark-top-only

Syntax [**no**] **qinq-mark-top-only**

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

Description When enabled (the encapsulation type of the access port where this SAP is defined as qinq), the **qinq-mark-top-only** command specifies which P-bits to mark during packet egress. When disabled, both set of P-bits are marked. When the enabled, only the P-bits in the top Q-tag are marked.

Default **no qinq-mark-top-only**

qos

Syntax **qos** *policy-id* [**shared-queuing** | **multipoint-shared**]
no qos

Context config>service>vprn>if>sap>egress
config>service>vprn>if>sap>ingress
config>service>vprn>sub-if>grp-if>sap>egress
config>service>vprn>sub-if>grp-if>sap>ingress

Description Associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP) or IP interface.

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the *policy-id* does not exist, an error will be returned.

The **qos** command is used to associate both ingress and egress QoS policies. The **qos** command only allows ingress policies to be associated on SAP or IP interface ingress and egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type will return an error.

When an ingress QoS policy is defined on an ingress IP interface that is bound to a VPRN, the policy becomes associated with every SAP on the VPRN and augments the QoS policy that is defined on each SAP. Packets that are bridged will be processed using the policy defined on the VPRN SAP; packets that are routed will be processed using the policy defined in the IES IP interface-binding context.

When an egress QoS policy is associated with an IP interface that has been bound to a VPRN, the policy becomes associated with every SAP on the VPRN and augments the egress QoS policy that is defined on each SAP. Packets that are bridged will be processed using the policy defined on the VPRN SAP; packets that are routed will be processed using the policy defined in the IP interface-binding context.

By default, no specific QoS policy is associated with the SAP or IP interface for ingress or egress, so the default QoS policy is used.

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

Parameters *policy-id* — 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 — Specify 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.

Values Multipoint or not present.

Default Present (the queue is created as non-multipoint).

scheduler-policy

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

Context config>service>vprn>if>sap>ingress
config>service>vprn>if>sap>egress
config>service>vprn>sub-if>grp-if>sap>egress
config>service>vprn>sub-if>grp-if>sap>ingress

Description This command applies an existing scheduler policy to an ingress or egress scheduler used by SAP 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.

multi-service-site

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

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

Description This command creates a new customer site or edits an existing customer site with the *customer-site-name* parameter. A customer site is an anchor point to create an ingress and egress virtual scheduler hierarchy. When a site is created, it must be assigned to a chassis slot or port with the exception of the 7750 SR-1 in which the slot is set to 1. When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).

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 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.

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** *queue-id*

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

Values 1 — 32

adaptation-rule

Syntax **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]
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** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir — The **cir** parameter defines the constraints enforced when adapting the CIR rate defined within the **queue** *queue-id* **rate** command. The **cir** parameter requires a qualifier that defines the

constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule — Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset.

Values **max** — The **max** (maximum) keyword is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) keyword is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

avg-frame-overhead

Syntax **avg-frame-overhead** *percent*
no avg-frame-overhead

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

Description This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a Sonet or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- Offered-load — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- Frame encapsulation overhead — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be 10000 x 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 x 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 x 1.1 or 550 octets.
- Frame based within-cir offered-load — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

  As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.
- Frame based PIR — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be 7500 x 1.1 or 8250 octets.
- Frame based within-pir offered-load — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

Port scheduler operation using frame transformed rates — The port scheduler uses the frame based rates to determine the maximum rates that each queue may receive during the within-cir and above-cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

SAP and subscriber SLA-profile average frame overhead override — The average frame overhead parameter on a sap-egress may be overridden at an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers an avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

Default **0**

Parameters *percent* — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

Values 0 — 100

cbs

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

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

Description This command can be used to override specific attributes of the specified queue's CBS parameters.

It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queue's CBS setting into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.

If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.

The **no** form of this command returns the CBS 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 or default

high-prio-only

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

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 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>vprn>if>sap>egress>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.

Values 0 — 131072 or default

mbs

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

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

Description This command can be used to override specific attributes of the specified queue's MBS parameters. The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not

force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The 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 returns the MBS size assigned to the queue to the value.

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

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

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

Description This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled 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 — 100000000

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 — 100000000, **max**, **sum**

Default 0

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.

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>service>vprn>if>sap>egress>sched-override>scheduler

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 then 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 ate 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

SAP Subscriber Management Commands

sub-sla-mgmt

Syntax [**no**] **sub-sla-mgmt**

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

Description This command entables the context to configure subscriber management parameters for this SAP.

Default **no sub-sla-mgmt**

def-sla-profile

Syntax **def-sla-profile** *default-sla-profile-name*
no def-sla-profile

Context config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt

Description This command specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sla-profile** context.

An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts. SLA profiles are maintained in two locations, the subscriber identification policy and the subscriber profile templates. After a subscriber host is associated with an SLA profile name, either the subscriber identification policy used to identify the subscriber or the subscriber profile associated with the subscriber host must contain an SLA profile with that name. If both the subscriber identification policy and the subscriber profile contain the SLA profile name, the SLA profile in the subscriber profile is used.

The **no** form of the command removes the default SLA profile from the SAP configuration.

Default **no def-sla-profile**

Parameters *default-sla-profile-name* — Specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sla-profile** context.

def-sub-profile

Syntax **def-sub-profile** *default-subscriber-profile-name*

Context config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt

Description This command specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-profile** context.

A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for

subscriber using the subscriber profile. Subscriber profiles also allow for specific SLA profile definitions when the default definitions from the subscriber identification policy must be overridden.

The **no** form of the command removes the default SLA profile from the SAP configuration.

Parameters *default-sub-profile* — Specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-profile** context.

multi-sub-sap

Syntax **multi-sub-sap** [*number-of-sub*]
no multi-sub-sap

Context config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt

Description This command configures the maximum number of subscribers for this SAP.

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

Parameters *number-of-sub* — Specifies the maximum number of subscribers for this SAP.

Values 2 — 8000

single-sub-parameters

Syntax **single-sub-parameters**

Context config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt

Description This command enables the context to configure single subscriber parameters for this SAP.

non-sub-traffic

Syntax **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*]
no non-sub-traffic

Context config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt>single-sub

Description This command configures non-subscriber traffic profiles.

The **no** form of the command removes removes the profiles and disables the feature.

Parameters **sub-profile** *sub-profile-name* — Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

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

subscriber *sub-ident-string* — Specifies an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

- For VPRN SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's sub-ident-string is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPRN destinations.

  If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

  If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.

  If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

  ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

profiled-traffic-only

Syntax [**no**] **profiled-traffic-only**

Context config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt>single-sub

Description This command enables profiled traffic only for this SAP.

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

sub-ident-policy

Syntax **sub-ident-policy** *sub-ident-policy-name*

Context config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt

Description This command associates a subscriber identification policy to this SAP. The subscriber identification policy must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-ident-policy** context.

Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.

For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet string. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.

When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.

The **no** form of the command removes the default subscriber identifcation policy from the SAP configuration.

Default **no sub-ident-policy**

Parameters *sub-ident-policy-name* — Specifies a subscriber identification policy for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-ident-policy** context.

srrp

Syntax [**no**] **srrp** *srrp-id*

Context config>service>vprn>sub-if>grp-if

Description This command creates an SRRP instance on a group IP interface. An SRRP instance manages all subscriber subnets within the group interfaces subscriber IP interface or other subscriber IP interfaces that are associated through a wholesale/retail relationship. Only one unique SRRP instance can be configured per group interface.

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

Default **no srrp**

Parameters *srrp-id* — Specifies a 32 bit instance ID that must be unique to the system. The instance ID must also match the instance ID used by the remote router that is participating in the same SRRP context. SRRP is intended to perform a function similar to VRRP where adjacent IP hosts within local subnets use a default gateway to access IP hosts on other subnets.

Values 1 — 4294967295

gw-mac

Syntax **gw-mac** *mac-address*
no gw-mac

Context config>service>vprn>sub-if>grp-if>srrp

Description This command overrides the default SRRP gateway MAC address used by the SRRP instance. Unless specified, the system uses the same base MAC address for all SRRP instances with the last octet overridden by the lower 8 bits of the SRRP instance ID. . The same SRRP gateway MAC address should be in-use by both the local and remote routers participating in the same SRRP context.

One reason to change the default SRRP gateway MAC address is if two SRRP instances sharing the same broadcast domain are using the same SRRP gateway MAC. The system will use the SRRP instance ID to separate the SRRP messages (by ignoring the messages that does not match the local

instance ID), but a unique SRRP gateway MAC is essential to separate the routed packets for each gateway IP address.

The **no** form of the command removes the explicit SRRP gateway MAC address from the SRRP instance. The SRRP gateway MAC address can only be changed or removed when the SRRP instance is shutdown.

Parameters *mac-address* — Specifies a MAC address that is used to override the default SRRP base MAC address

Values Any MAC address except all zeros, broadcast or multicast addresses. The offset is expressed in normal Ethernet MAC address notation. The defined gw-mac cannot be 00:00:00:00:00:00, ff:ff:ff:ff:ff:ff or any multicast address.

If not specified, the system uses the default SRRP gateway MAC address with the last octet set to the 8 least significant bits of the SRRP instance ID.

keep-alive-interval

Syntax **keep-alive-interval** *interval*
no keep-alive-interval

Context config>service>vprn>sub-if>grp-if>srrp

Description This command defines the interval between SRRP advertisement messages sent when operating in the master state. The interval is also the basis for setting the master-down timer used to determine when the master is no longer sending. The system uses three times the keep-alive interval to set the timer. Every time an SRRP advertisement is seen that is better then the local priority, the timer is reset. If the timer expires, the SRRP instance assumes that a master does not exist and initiates the attempt to become master.

When in backup state, the SRRP instance takes the keep-alive interval of the master as represented in the masters SRRP advertisement message. Once in master state, the SRRP instance uses its own configured keep-alive interval.

The keep-alive-interval may be changed at anytime, but will have no effect until the SRRP instance is in the master state.

The **no** form of the command restores the default interval.

Parameters *interval* — Specifies the interval between SRRP advertisement messages sent when operating in the master state.

Values 1 — 100 hundreds of milli-seconds

Default 1

message-path

Syntax **message-path** *sap-id*
no message-path

Context config>service>vprn>sub-if>grp-if>srrp

Description This command defines a specific SAP for SRRP in-band messaging. A message-path SAP must be defined prior to activating the SRRP instance. The defined SAP must exist on the SRRP instances group IP interface for the command to succeed and cannot currently be associated with any dynamic or static subscriber hosts. Once a group IP interface SAP has been defined as the transmission path for SRRP Advertisement messages, it cannot be administratively shutdown, will not support static or dynamic subscriber hosts and cannot be removed from the group IP interface.

The SRRP instance message-path command may be executed at anytime on the SRRP instance. Changing the message SAP will fail if a dynamic or static subscriber host is associated with the new SAP. Once successfully changed, the SRRP instance will immediately disable anti-spoof on the SAP and start sending SRRP Advertisement messages if the SRRP instance is activated.

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

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

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

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

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

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b

	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

policy

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

Context config>service>vprn>sub-if>grp-if>srrp

Description This command associates one or more VRRP policies with the SRRP instance. A VRRP policy is a collection of connectivity and verification tests used to manipulate the in-use priorities of VRRP and SRRP instances. A VRRP policy can test the link state of ports, ping IP hosts, discover the existence of routes in the routing table or the ability to reach L2 hosts. When one or more of these tests fail, the VRRP policy has the option of decrementing or setting an explicit value for the in-use priority of an SRRP instance.

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

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

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

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

Values 1 — 9999

priority

Syntax **priority** *priority*
no priority

Context config>service>vprn>sub-if>grp-if>srrp

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

The base priority of an SRRP instance can be managed by VRRP policies. A VRRP policy defines a set of connectivity or verification tests which, when they fail, may lower an SRRP instances base priority (creating an in-use priority for the instance). Every time an SRRP instances in-use priority changes when in master state, it sends an SRRP advertisement message with the new priority. If the dynamic priority drops to zero or receives an SRRP Advertisement message with a better priority, the SRRP instance transitions to the *becoming backup* state.

When the priority command is not specified, or the no priority command is executed, the system uses a default base priority of 100. The priority command may be executed at anytime.

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

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

Values 1 — 254

Default 100

Interface VRRP Commands

vrrp

Syntax **vrrp** *virtual-router-id* [**owner**]
no vrrp *virtual-router-id*

Context config>service>vprn>interface

Description This command creates or edits a Virtual Router ID (VRID) on the service IP interface. A VRID is internally represented in conjunction with the IP interface name. This allows the VRID to be used on multiple IP interfaces while representing different virtual router instances.

Two VRRP nodes can be defined on an IP interface. One, both, or none may be defined as owner. The nodal context of **vrrp** *virtual-router-id* is used to define the configuration parameters for the VRID.

The **no** form of this command removes the specified VRID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the VRID. The VRID does not need to be shutdown in order to remove the virtual router instance.

Default **No default**

Parameters *virtual-router-id* — The virtual-router-id parameter specifies a new virtual router ID or one that can be modified on the IP interface.

Values 1 — 255

authentication-key

Syntax **authentication-key** [*authentication-key* **|** *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>service>vprn>if>vrrp

Description The **authentication-key** command, within the **vrrp** *virtual-router-id* context, is used to assign a simple text password authentication key to generate master VRRP advertisement messages and validate received VRRP advertisement messages.

The authentication-key command is one of the few commands not affected by the presence of the owner keyword. If simple text password authentication is not required, the authenticaton-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, altering the simple text password used when authentication-type password authentication method is used by the virtual router instance. The authentication-type password command does not need to be executed prior to defining the authentication-key command.

To change the current in-use password key on multiple virtual router instances:

- Identify the current master
- Shutdown the virtual router instance on all backups
- Execute the authentication-key command on the master to change the password key

- Execute the authentication-key command and no shutdown command on each backup key

The **no** form of this command restores the default null string to the value of key.

Default No default. The authentication data field contains the value 0 in all 16 octets.

Parameters *authentication-key* — 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.

authentication-type

Syntax **authentication-type** {**password** | **message-digest**}
no authentication-type

Context config>service>vprn>if>vrrp

Description The **authentication-type** command, within the **vrrp** *virtual-router-id* context, is used to assign the authentication method to generate master VRRP advertisement messages and validate received VRRP advertisement messages.

NOTE: The authentication management for VRRP closely follows the authentication management format used for IS-IS.

The **authentication-type** command is one of the commands not affected by the presence of the owner keyword. If authentication is not required, the authenticaton-type command must not be executed. If the command is re-executed with a different authentication type defined, the new type will be used. If the no authentication-type command is executed, authentication is removed and no authentication is performed. The authentication-type command may be executed at any time, altering the authentication method used by the virtual router instance.

The **no** form of this command removes authentication from the virtual router instance. All VRRP Advertisement messages sent will have the Authentication Type field set to 0 and the Authentication Data fields will contain 0 in all octets. VRRP Advertisement messages received with Authentication Type fields containing a value other than 0 will be discarded.

password — The password keyword identifies VRRP Authentication Type 1. Type 1 requires the definition of a string of eight octets long using the authentication-key command. All transmitted VRRP Advertisement messages must have the Authentication Type field set to 1 and the Authentication Data fields must contain the authentication-key password.

All received VRRP advertisement messages must contain a value of 1 in the Authentication Type field and the Authentication Data fields must match the defined authentication-key. All other received messages will be silently discarded.

message-digest — The message-digest keyword identifies VRRP Authentication Type 2. Type 2 defines a lower IP layer MD5 authentication mechanism using HMAC and IP authentication header standards. An MD5 key must be defined using the message-digest-key command. All transmitted VRRP advertisement messages must have the Authentication Type field set to 2 and the Authentication Data fields must contain 0 in all octets. The message-digest key is used in the hashing process when populating the IP Authentication Header fields. A sequential incrementing counter (set to zero when the message-digest-key is set) is incremented and then used in the IP Authentication Header to prevent replay attacks on authorized participating virtual router instances.

All received VRRP advertisement messages must contain a value of 2 in the Authentication Type field and the Authentication Data fields are ignored. The message must have been authorized by the lower layer IP Authentication Header process with the sequential counter field and the source IP address presented to the virtual router instance. To track the validity of the received counter, the virtual router instance maintains a master counter table containing up to 32 source IP addresses and the last received counter value. Populate the table as follows:

1. Check to see if source IP address exists in table.
   - If non-existent, create an entry if available.
   - If no entry is available, delete the oldest and create an entry.
     The new entry should have a counter value of zero.
2. Compare the message counter value to the entry value (0 if new entry or equal to the previous message counter from the source IP address).
   - If the message counter is not greater than the entry counter value, silently discard the packet.
   - If the message counter is greater than the entry counter value, accept the message for further checking and replace the entry counter value with the message counter value and time stamp the entry.

backup

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

Context config>service>vprn>if>vrrp

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

init-delay

Syntax **init-delay** *seconds*
no init-delay

Context config>service>vprn>if>vrrp

Description This command configures a VRRP initialization delay timer.

Default **no init-delay**

Parameters *seconds* — Specifies the initialization delay timer for VRRP, in seconds.

Values 1 — 65535

mac

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

Context config>service>VPRN>if>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 *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.

master-int-inherit

Syntax [**no**] **master-int-inherit**

Context config>service>vprn>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>vprn>interface

Description This command sets the advertisement timer and indirectly sets the master down timer on the virtual router instance. The message-interval setting must be the same for all virtual routers participating as a virtual router. Any VRRP advertisement message received with an Advertisement Interval field different than the virtual router instance configured message-interval value will be silently discarded.

The message-interval command is available in both non-owner and owner **vrrp** *virtual-router-id* nodal contexts. If the message-interval command is not executed, the default message interval of 1 second will be used.

The **no** form of this command restores the default message interval value of 1 second to the virtual router instance.

Parameters *seconds* — The number of seconds that will transpire before the advertisement timer expires.

Values 1 — 255

Default 1

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

Values 100 — 900

ping-reply

Syntax [**no**] **ping-reply**

Context config>service>vprn>if>vrrp

Description This command enables the non-owner master to reply to ICMP Echo Requests directed at the virtual router instances IP addresses. The ping request can be received on any routed interface.

Ping must not have been disabled at the management security level (either on the parental IP interface or based on the Ping source host address). When ping-reply is not enabled, ICMP Echo Requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to ICMP Echo Requests regardless of the setting of ping-reply configuration.

The ping-reply command is only available in non-owner **vrrp** *virtual-router-id* nodal context. If the ping-reply command is not executed, ICMP Echo Requests to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all ICMP Echo Request messages destined to the non-owner virtual router instance IP addresses.

Default no ping-reply

policy

Syntax **policy** *vrrp-policy-id*
no policy

Context config>service>vprn>if>vrrp

Description This command associates a VRRP priority control policy with the virtual router instance (non-owner context only).

Parameters *vrrp-policy-id* — Specifies a VRRP priority control policy.

Values 1 — 9999

preempt

Syntax **preempt**
no preempt

Context config>service>vprn>interface

Description This command provides the ability of overriding an existing non-owner master to the virtual router instance. Enabling preempt mode is recommended 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** *priority*
no priority

Context config>service>vprn>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

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

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

Description This command enables the non-owner master to reply to TCP port 23 Telnet Requests directed at the virtual router instance's IP addresses. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Proper login and CLI command authentication is still enforced.

When telnet-reply is not enabled, TCP port 23 Telnet packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to Telnet Requests regardless of the telnet-reply configuration.

The telnet-reply command is only available in non-owner **VRRP** nodal context. If the telnet-reply command is not executed, Telnet packets to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all Telnet packets destined to the non-owner virtual router instance IP addresses.

Default **no telnet-reply**

traceroute-reply

Syntax [**no**] **traceroute-reply**

Context config>service>vprn>if>vrrp

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 This command configures a Protocol Independent Multicast (PIM) instance in the VPRN service. When an PIM instance is created, the protocol is enabled.

The **no** form of the command deletes the PIM protocol instance removing all associated configuration parameters.

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 7750 SR OS supports PIM sparse mode (PIM-SM).

Default **none**

apply-to

Syntax **apply-to** {**all** | **none**}

Context config>service>vprn>pim

Description This command creates a PIM interface with default parameters.

If a manually created interface or modified interface is deleted, the interface will be recreated when the **apply-to** command is executed. If PIM is not required on a specific interface, then execute a **shutdown** command.

The **apply-to** command is saved first in the PIM configuration structure, all subsequent commands either create new structures or modify the defaults as created by the **apply-to** command.

Default **none** (keyword)

Parameters **all —** Specifies that all VPRN and non-VPRN interfaces are automatically applied in PIM.

none — No interfaces are automatically applied in PIM. PIM interfaces must be manually configured.

import

Syntax **import** {**join-policy** | **register-policy**} [*policy-name* [.. *policy-name*] *policy-name*]
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.

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.

interface

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

Context config>service>vprn>pim

Description This command enables PIM on an interface and enables the context to configure interface-specific parameters.

By default interfaces are activated in PIM based on the **apply-to** command, and do not have to be configured on an individual basis unless the default values must be changed.

The **no** form of the command deletes the PIM interface configuration for this interface. If the **apply-to** command parameter is configured, then the **no interface** form must be saved in the configuration to avoid automatic (re)creation after the next **apply-to** is executed as part of a reboot.

The **shutdown** command can be used to disable an interface without removing the configuration for the interface.

Default Interfaces are activated in PIM based on the apply-to command.

Parameters *ip-int-name* — Specify the interface name. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

bfd-enable

Syntax [**no**] **bfd-enable**

Context config>service>vprn>pim>interface

Description This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the

protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

The **no** form of this command removes BFD from the associated IGP protocol adjacency.

Default **no bfd-enable**

bsm-check-rtr-alert

Syntax [**no**] **bsm-check-rtr-alert**

Context config>service>vprn>pim>interface

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>interface
config>service>vprn>pim>mdt>default

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>interface
config>service>vprn>pim>mdt>default

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:

(hello-interval * 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>interface
config>service>vprn>pim>mdt>default

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**

max-groups

Syntax **max-groups** *value*
no max-groups

Context config>service>vprn>pim>interface

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>interface

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.

priority

Syntax **priority** *dr-priority*
no priority

Context config>service>vprn>pim>interface

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 *dr-priority* — 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** *dr-priority*]
no sticky-dr

Context config>service>vprn>pim>interface

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 designateed router (DR) will be modified to the value configured in *dr-priority.* This is done to avoid the delays in forwarding caused by DR recovery, when switching back to the old DR on a LAN when it comes back up.

By enabling **sticky-dr** 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** *dr-priority* **—** Sets the DR priority to be sent in PIM Hello messages following the election of that interface as the DR, when sticky-dr operation is enabled.

Values 1 — 4294967295

three-way-hello

Syntax **three-way-hello** [**compatibility-mode**]

no three-way-hello

Context config>service>vprn>pim>interface
config>service>vprn>pim>mdt>default

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>interface
config>service>vprn>pim>mdt>default

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**

mdt

Syntax **mdt**

Context config>service>vprn>pim>interface

Description This command enables the context for a multicast distribution tree (MDT) to carry multicast traffic from customer sites associated with the multicast domain. Multicast-capable routers create distribution trees that control the path that IP multicast traffic takes through the network to deliver traffic to all receivers. There are two types of MDTs, source trees and shared trees. The root of the source tree is the source of the multicast tree whose branches form a spanning tree through the network to the receivers. It is also referred to as a shortest path tree (SPT) because the tree uses the shortest path through the network.

Shared trees use a common root that is placed at a specific place in the network. This shared root is called the rendezvous point (RP).

All PEs that are configured with the same MDT address will become members of this group and receive multicast traffic from each other.

The source address used in MDT group address packets is the loopback address configured for the VPRN, if the loopback address is removed the service we will attempt to find another loop-back address for the VPRN instance, if no loopback address exists then multicast tunnel for the VPRN instance will be administratively down. The show command will reflect the reason why the PIM-SM instance is down.

Addressing conflicts in the core can be avoided by installing import policies on the main PIM access interfaces.

To enable multicast in a VPRN this parameter must be configured. If it is not configured, no PIM-SM will not be initialized for this VPRN, and the show command will indicate that the default MDT address is missing. If the address is removed using the no form of this command, Multicast will be shut down for this instance and an error indication is displayed when a show command is executed.

Use the **no** form of this command to remove default MDT address from the configuration.

Default **none**

data

Syntax **data** {*grp-ip-address/mask* | *grp-ip-address netmask*}

Context config>service>vprn>pim>mdt

Description This command configures a pool of addresses that can be used to generate data only MDT tunnels.

Parameters *grp-ip-address* — The multicast group IP address expressed in dotted decimal notation.

Values 224.0.0.0 — 239.255.255.255

mask — The mask associated with the IP prefix expressed as a mask length or in dotted decimal notation; for example /16 for a sixteen-bit mask. The mask can also be entered in dotted decmial notation (255.255.0.0).

Values 4 — 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)

data-delay-interval

Syntax **data-delay-interval** *value*
no data-delay-interval

Context config>service>vprn>pim>mdt

Description This command specifies the interval, in seconds, before the provider edge (PE) router connected to the source switches traffic from default Multicast Distribution Tree (MDT) to the data MDT group.

Default **3 seconds**

Parameters *value* — Specifies the data delay interval in seconds.

Values 3 — 180

data-threshold

Syntax **data-threshold** {*c-grp-ip-address/mask* | *c-grp-ip-address netmask*} *mdt-threshold*

Context config>service>vprn>pim>mdt

Description This command configures the threshold for a group prefix.

Parameters *grp-ip-address* — The multicast group IP address expressed in dotted decimal notation.

Values 224.0.0.0 — 239.255.255.255

mask — The mask associated with the IP prefix expressed as a mask length or in dotted decimal notation; for example /16 for a sixteen-bit mask. The mask can also be entered in dotted decmial notation (255.255.0.0).

Values 4 — 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)

mdt-threshold — Specifies the threshold, in kilo-bits per second (kbps), for the group to which this C-(S,G) belongs. For a C-group G configured with a threshold, a C-(S,G) is mapped to a Data Multicast Tunnel (MT) only if the C-(S,G)'s rate exceeds this configured threshold.

default

Syntax **default** *grp-ip-address*
no default

Context config>service>vprn>pim>mdt

Description This command configures a default multicast distribution tree (MDT) group address used by the core instance of PIM to identify multicast traffic for this VPRN instance. All PE's that are configured with the same MDT address will become members of this group and receive multicast traffic from each other.

The **no** form of this command removes the MDT default address from the configuration.

Parameters *grp-ip-address* — The multicast IP address for the group.

Values 224.0.1.0 — 239.255.255.255

join-tlv-packing-disable

Syntax [**no**] **join-tlv-packing-disable**

Context config>service>vprn>pim>mdt

Description This command specifies enables the packing of MDT join TLVs. If multiple Join TLVs are available at the time they are transmitted. The TLVs are packed into a single UDP PDU instead of sending separate UDP PDUs. In scaling scenarios, this packing makes more efficient use of packet buffers and helps with better convergence.

non-dr-attract-traffic

Syntax [**no**] **non-dr-attract-traffic**

Context config>service>vprn>pim

Description This command specifies whether the router should ignore the designated router state and attract traffic even when it is not the designater router.

An operator can configure an interface (router or IES or VPRN interfaces) to IGMP and PIM. The interface IGMP state will be synchronized to the backup node if it is associated with the redundant peer port. The interface can be configured to use PIM which will cause multicast streams to be sent to the elected DR only. The DR will also be the router sending traffic to the DSLAM. Since it may be required to attract traffic to both routers a flag non-dr-attract-trafffic can be used in the PIM context to have the router ignore the DR state and attract traffic when not DR. Note that while using this flag the router may not send the stream down to the DSLAM while not DR.

When enabled, the designated router state is ignored. When disabled, **no non-dr-attract-traffic**, the designated router value is honored.

Default **no non-dr-attract-traffic**

rp

Syntax **rp**

Context config>service>vprn>pim

Description This command enables access to the context to configure the rendezvous point (RP)) of a PIM protocol instance.

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.

Default **rp enabled when PIM is enabled.**

anycast

Syntax [**no**] **anycast** *rp-ip-address*

Context config>service>vprn>pim>rp

Description This command configures a PIM anycast protocol instance for the RP being configured. Anycast enables fast convergence when a PIM RP router fails by allowing receivers and sources to rendezvous at the closest RP.

The **no** form of the command removes the anycast instance from the configuration.

Default **none**

Parameters *rp-ip-address* — Configure the loopback IP address shared by all routes that form the RP set for this anycast instance. Only a single address can be configured. If another anycast command is entered with an address then the old address will be replaced with the new address. If no ip-address is entered then the command is simply used to enter the anycast CLI level.

Values Any valid loopback address configured on the node.

rp-set-peer

Syntax [**no**] **rp-set-peer** *ip-address*

Context config>service>vprn>pim>rp>anycast

Description This command configures a peer in the anycast rp-set. The address identifies the address used by the other node as the RP 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**

Values Any valid ip-address within the scope outlined above.

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

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

hash-mask-len

Syntax **hash-mask-len** *hash-mask-length*
no hash-mask-len

Context config>service>vprn>pim>rp>bsr-candidate

Description This command is used to configure the length of a mask that is to be combined with the group address before the hash function is called. All groups with the same hash map to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This mechanism is used to map one group or multiple groups to an RP.

Parameters *hash-mask-length* — The hash mask length.

Values 0 — 32

priority

Syntax **priority** *bootstrap-priority*

Context config>service>vprn>pim>rp>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 *bootstrap-priority* — The priority to become the bootstrap router.

Values 0 — 255

Default 0 (the router is not eligible to be the bootstrap router)

rp-candidate

Syntax **rp-candidate**

Context config>service>vprn>pim>rp

Description This command enables the context to configure the candidate rendezvous point (RP) parameters.

Default **Enabled when PIM is enabled.**

group-range

Syntax [**no**] **group-range** {*grp-ip-address*/*mask* | *grp-ip-address* [*netmask*]}

Context config>service>vprn>pim>rp>rp-candidate
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 No group addresses/masks are configured.

Parameters *group-ip-address* — Specify the addresses or address ranges that this router can be an RP.

mask — Specify the address mask with the address to define a range of addresses.

netmask — Specify 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)

holdtime

Syntax **holdtime** *holdtime*
no holdtime *holdtime*

Context config>service>vprn>pim>rp>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.

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 dyanamically learned Rendezvous Point (RP).

When enabled, the static group-to-RP mappings take precedence over the dynamically learned mappings.

Default **no override**

spt-switchover-threshold

Syntax **spt-switchover-threshold** {*grp-ip-address/mask* | *grp-ip-address netmask*} *spt-threshold*
no spt-switchover-threshold {*grp-ip-address/mask* | *grp-ipaddress netmask*}

Context config>service>vprn>pim

Description This command configures a shortest path tree (SPT tree) switchover threshold for a group prefix.

Parameters *grp-ip-address* — Specify the multicast group 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)

spt-threshold — Specifies the configured threshold in kilo-bits per second(kbps) for the group to which this (S,G) belongs. For a group G configured with a threshold, switchover to SPT for an (S,G) is attempted only if the (S,G)'s rate exceeds this configured threshold.

ssm-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**

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**

advertise-inactive

Syntax [**no**] **advertise-inactive**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command enables or disables the advertising of inactive BGP routers to other BGP peers.

By default, BGP only advertises BGP routes to other BGP peers if a given BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the most preferred route within the system for a given destination.

Default **no advertise-inactive**

aggregator-id-zero

Syntax [**no**] **aggregator-id-zero**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command is used to set the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes that contain different AS paths.

When BGP is aggregating routes, it adds the aggregator path attribute to the BGP update messages. By default, BGP adds the AS number and router ID to the aggregator path attribute.

When this command is enabled, BGP adds the router ID to the aggregator path attribute. This command is used at the group level to revert to the value defined under the global level, while this command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of the command used at the global level reverts to default where BGP adds the AS number and router ID to the aggregator path attribute.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no aggregator-id-zero** — 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

Context config>service>vprn>bgp

Description This command specifies how the Multi-Exit Discriminator (MED) path attribute is used in the BGP route selection process. The MED attribute is always used in the route selection process regardless of the peer AS that advertised the route. This parameter determines what MED value is inserted in the RIB-IN.

If this parameter is not configured, only the MEDs of routes that have the same peer ASs are compared.

The **no** form of the command removes the parameter from the configuration.

Default **no always-compare-med** — Only compare MEDs of routes that have the same peer AS.

Parameters **zero —** Specifies that for routes learned without a MED attribute that a zero (0) value is used in the MED comparison. The routes with the lowest metric are the most preferred.

infinity — Specifies for routes learned without a MED attribute that a value of infinity (4294967295) is used in the MED comparison. This in effect makes these routes the least desirable.

as-path-ignore

Syntax [**no**] **as-path-ignore**

Context config>service>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**

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** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures the BGP authentication key.

Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD-5 message-based digest. The authentication key can be any combination of letters or numbers from 1 to 16.

The no form of the command removes the authentication password from the configuration and effectively disables authentication.

Default Authentication is disabled and the authentication password is empty.

Parameters *authentication-key* — The authentication key. The key can be any combination of ASCII characters up to 255 characters in length (unencrypted). If spaces are used in the string, 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.

auth-keychain

Syntax **auth-keychain** *name*

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 *name* — Specifies the name of an existing keychain, up to 32 characters, to use for the specified TCP session or sessions.

cluster

Syntax **cluster** *cluster-id*
no cluster

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures the cluster ID for a route reflector server.

Route reflectors are used to reduce the number of IBGP sessions required within an AS. Normally, all BGP speakers within an AS must have a BGP peering with every other BGP speaker in an AS. A route reflector and its clients form a cluster. Peers that are not part of the cluster are considered to be non-clients.

When a route reflector receives a route, first it must select the best path from all the paths received. If the route was received from a non-client peer, then the route reflector sends the route to all clients in the cluster. If the route came from a client peer, the route reflector sends the route to all non-client peers and to all client peers except the originator.

For redundancy, a cluster can have multiple route reflectors.

Confederations can also be used to remove the full IBGP mesh requirement within an AS.

The **no** form of the command deletes the cluster ID and effectively disables the Route Reflection for the given group.

Default **no cluster** — No cluster ID is defined.

Parameters *cluster-id* — The route reflector cluster ID is expressed in dot decimal notation.

Values Any 32 bit number in dot decimal notation. (0.0.0.1 — 255.255.255.255)

connect-retry

Syntax **connect-retry** *seconds*
no connect-retry

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures the BGP connect retry timer value in seconds.

When this timer expires, BGP tries to reconnect to the configured peer. This configuration parameter can be set at three levels: global level (applies to all peers), peer-group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to the default value.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **120** seconds

Parameters *seconds* — The BGP Connect Retry timer value in seconds, expressed as a decimal integer.

Values 1 — 65535

damping

Syntax [**no**] **damping**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command enables BGP 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-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.

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**

export

Syntax **export** *policy* [*policy...*]
no export

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command specifies the export policies to be used to control routes advertised to 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 *policy* — A route policy statement name.

group

Syntax **group** *name*
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.

neighbor

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

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

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 *ip-address* — The IP address of the BGP peer router in dotted decimal notation.

hold-time

Syntax **hold-time** *seconds*
no hold-time

Context config>service>vprn>bgp
config>service>vprn>bgp>group

Description This command configures the BGP hold time, expressed in seconds.

The BGP hold time specifies the maximum time BGP waits between successive messages (either keepalive or update) from its peer, before closing the connection. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

Even though the router OS implementation allows setting the **keepalive** time separately, the configured **keepalive** timer is overridden by the **hold-time** value under the following circumstances:

1. If the specified hold-time is less than the configured **keepalive** time, then the operational **keepalive** time is set to a third of the **hold-time**; the configured **keepalive** time is not changed.
2. If the **hold-time** is set to zero, then the operational value of the **keepalive** time is set to zero; the configured **keepalive** time is not changed. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of the command used at the global level reverts to the default value.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **90** seconds

Parameters *seconds* — The hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is up permanently.

Values 0, 3 — 65535

ibgp-multipath

Syntax [**no**] **ibgp-multipath**

Context config>service>vprn>bgp

Description This command defines the type of IBGP multipath to use when adding BGP routes to the route table if the route resolving the BGP nexthop offers multiple nexthops.

The **no** form of the command resets the IBGP multipath feature to the default value (*load-balance*).

Default **ibgp-multipath load-balance**

import

Syntax **import** *policy* [*policy*...]
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 *policy* — A route policy statement name.

keepalive

Syntax **keepalive** *seconds*
no keepalive

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures the BGP keepalive timer. A keepalive message is sent every time this timer expires.

The **keepalive** parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in 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

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, 7750 SR OS uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGP peers. This command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of the command removes the configured local-address for BGP.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

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 local-as

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures a BGP virtual autonomous system (AS) number.

In addition to the AS number configured for BGP in the `config>router>autonomous-system` context, a virtual (local) AS number is configured.The virtual AS number is added to the as-path message before the router's AS number makes the virtual AS the second AS in the as-path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). Thus, by specifying this at each neighbor level, it is possible to have a separate as-number per EBGP session.

When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The **private** attribute can be added or removed dynamically by reissuing the command.

Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local AS number. Changing the local AS at the global level in an active BGP instance causes BGP to 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 *as-number* — 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.

local-preference

Syntax **local-preference** *local-preference*
no local-preference

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command enables setting the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute.

This value is used if the BGP route arrives from a BGP peer without the **local-preference** integer set.

The specified value can be overridden by any value set via a route policy. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command at the global level specifies that incoming routes with local-preference set are not overridden and routes arriving without local-preference set are interpreted as if the route had local-preference value of 100.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no local-preference** — Does not override the local-preference value set in arriving routes and analyze routes without local preference with value of 100.

Parameters *local-preference* — 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 This command configures how the BGP peer session handles loop detection in the AS path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

Note that dynamic configuration changes of **loop-detect** are not recognized.

The **no** form of the command used at the global level reverts to default, which is **loop-detect ignore-loop**.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **loop-detect ignore-loop**

Parameters **drop-peer —** Sends a notification to the remote peer and drops the session.

discard-route — Discards routes received with loops in the AS path.

ignore-loop — Ignores routes with loops in the AS path but maintains peering.

off — Disables loop detection.

med-out

Syntax **med-out** {**number | igp-cost**}
no med-out

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set.

The specified value can be overridden by any value set via a route policy.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to default where the MED is not advertised.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no med-out**

Parameters *number* — 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-as-origination

Syntax **min-as-origination** *seconds*
no min-as-origination

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures the minimum interval, in seconds, at which a path attribute, originated by the local router, can be advertised to a peer.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to default.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **15** seconds

Parameters *seconds* — The minimum path attribute advertising interval in seconds, expressed as a decimal integer.

Values 2 — 255

min-route-advertisement

Syntax **min-route-advertisement** *seconds*
no min-route-advertisement

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

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 *seconds* — The minimum route advertising interval, in seconds, expressed as a decimal integer.

Values 2 — 255

multihop

Syntax **multihop** *ttl-value*
no multihop

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures the time to live (TTL) value entered in the IP header of packets sent to an EBGP peer multiple hops away.

This parameter is meaningful only when configuring EBGP peers. It is ignored if set for an IBGP peer.

The **no** form of the command is used to convey to the BGP instance that the EBGP peers are directly connected.
The **no** form of the command reverts to default values.

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** *integer*
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

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** *as-number*

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.

For EBGP peers, the peer AS number configured must be different from the autonomous system number configured for this router under the global level since the peer will be in a different autonomous system than this router

For IBGP peers, the peer AS number must be the same as the autonomous system number of this router configured under the global level.

This is a required command for each configured peer. This may be configured under the group level for all neighbors in a particular group.

Default No AS numbers are defined.

Parameters *as-number* — The autonomous system number, expressed as a decimal integer.

Values 1 — 65535

preference

Syntax [**no**] **preference** *preference*

Context config>service>vprn>bgp
config>service>vprn>bgp>group

Description This command configures the route preference for routes learned from the configured peer(s).

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The lower the preference the higher the chance of the route being the active route. The OS assigns BGP routes highest default preference compared to routes that are direct, static or learned via MPLS or OSPF.

The **no** form of the command used at the global level reverts to default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **170**

Parameters *preference* — The route preference, expressed as a decimal integer.

Values 1 — 255

prefix-limit

Syntax **prefix-limit** *limit*
no prefix-limit

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures the maximum number of routes BGP can learn from a peer.

When the number of routes reaches 90% of this limit, an SNMP trap is sent. When the limit is exceeded, the BGP peering is dropped and disabled.

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

Default **no prefix-limit**

Parameters *limit* — The number of routes that can be learned from a peer, expressed as a decimal integer.

Values 1 — 4294967295

remove-private

Syntax [**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.

type

Syntax [**no**] **type** {**internal** | **external**}

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command designates the BGP peer as type internal or external.

The type of **internal** indicates the peer is an IBGP peer while the type of external indicates that the peer is an EBGP peer.

By default, the OS derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered **internal**. If the local AS is different, then the peer is considered **external**.

The **no** form of the command used at the group level reverts to the default value.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no type** — Type of neighbor is derived on the local AS specified.

Parameters **internal —** Configures the peer as internal.

external — Configures the peer as external.

ttl-security

Syntax **ttl-security** *min-ttl-value*
no ttl-security

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description Configure TTL security parameters for incoming packets.

Parameters *min-ttl-value* — Specify the minimum TTL value for an incoming BGP packet.

Values 1 — 255

Default 1

OSPF Commands

ospf

Syntax [**no**] **ospf**

Context config>service>vprn

Description This command enables access to the context to enable an OSPF protocol instance.

When an OSPF instance is created, the protocol is enabled. To start or suspend execution of the OSPF protocol without affecting the configuration, use the **no shutdown** command.

The **no** form of the command deletes the OSPF protocol instance removing all associated configuration parameters.

Default **no ospf** — The OSPF protocol is not enabled.

area

Syntax [**no**] **area** *area-id*

Context config>service>vprn>ospf

Description This command creates the context to configure an OSPF area. An area is a collection of network segments within an AS that have been administratively grouped together. The area ID can be specified in dotted decimal notation or as a 32-bit decimal integer.

The **no** form of the command deletes the specified area from the configuration. Deleting the area also removes the OSPF configuration of all the interfaces, virtual-links, and address-ranges etc., that are currently assigned to this area.

Default **no area** — No OSPF areas are defined.

Parameters *area-id* — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

Values 0.0.0.0 — 255.255.255.255 (dotted decimal), 0 — 4294967295 (decimal integer)

area-range

Syntax **area-range** *ip-prefixImask* [**advertise** | **not-advertise**]
no area-range *ip-prefixImask*

Context config>service>vprn>ospf>area
ospf>service>vprn>nssa

Description This command creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression. When a range is created, the range is configured to be advertised or not advertised into other areas. Multiple range commands may be used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.

ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.

The **no** form of the command deletes the range (non) advertisement.

Default **no area-range** — No range of addresses are defined.

Special Cases **NSSA Context —** In the NSSA context, the option specifies that the range applies to external routes (via type-7 LSAs) learned within the NSSA when the routes are advertised to other areas as type-5 LSAs.

Area Context — If this command is not entered under the NSSA context, the range applies to summary LSAs even if the area is an NSSA.

Parameters *ip-prefix* — The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.

Values 0.0.0.0 - 255.255.255.255

mask — The subnet mask for the range expressed as a decimal integer mask length or in dotted decimal notation.

Values 0 — 32 (mask length), 0.0.0.0 — 255.255.255.255 (dotted decimal)

advertise | **not-advertise —** Specifies whether or not to advertise the summarized range of addresses into other areas. The **advertise** keyword indicates the range will be advertised, and the keyword **not-advertise** indicates the range will not be advertised.
The default is **advertise**.

blackhole-aggregate

Syntax [**no**] **blackhole-aggregate**

Context config>service>vprn>ospf>area

Description This command installs a low priority blackhole route for the entire aggregate. Exisiting 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.

Use the no form of this command to remove this option.

Default **blackhole-aggregate**

interface

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

Context config>service>vprn>ospf>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.

advertise-subnet

Syntax [**no**] **advertise-subnet**

Context config>service>vprn>ospf>area>interface

Description This command enables advertising point-to-point interfaces as subnet routes (network number and mask). When disabled, point-to-point interfaces are advertised as host routes.

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-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-link

Description This command configures the password used by the OSPF interface or virtual-link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.

All neighboring routers must use the same type of authentication and password for proper protocol communication. If the **authentication-type** is configured as password, then this key must be configured.

By default, no authentication key is configured.

The **no** form of the command removes the authentication key.

Default **no authentication-key** - No authentication key is defined.

Parameters *authentication-key* — The authentication key. The key can be any combination of ASCII characters up to 8 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

hash-key — The hash key. The key can be any combination of ASCII characters up to 22 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

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

authentication-type

Syntax **authentication-type** {**password** | **message-digest**}
no authentication-type

Context config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-link

Description This command enables authentication and specifies the type of authentication to be used on the OSPF interface.

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.

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

dead-interval

Syntax **dead-interval** *seconds*
no dead-interval

Context config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-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** seconds

Special Cases **OSPF Interface —** If the **dead-interval** configured applies to an interface, then all nodes on the subnet must have the same dead interval.

Virtual Link — If the **dead-interval** configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same dead interval.

Parameters *seconds* — The dead interval expressed as a decimal integer.

Values 2 — 2147483647 seconds

hello-interval

Syntax **hello-interval** *seconds*
no hello-interval

Context config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-link

Description This command configures the interval between OSPF hellos issued on the interface or virtual 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.

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>interface

Description This command configures the interface type to be either broadcast or point-to-point.

Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the broadcast adjacency maintenance overhead if the Ethernet link provided the link is used as a point-to-point.

If the interface type is not known at the time the interface is added to OSPF and subsequently the IP interface is bound (or moved) to a different interface type, this command must be entered manually.

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

Default **point-to-point** — If the physical interface is SONET.

broadcast — If the physical interface is Ethernet or unknown.

Special Cases **Virtual-Link —** A virtual link is always regarded as a point-to-point interface and not configurable.

Parameters **broadcast —** Configures the interface to maintain this link as a broadcast network. To significantly improve adjacency forming and network convergence, a network should be configured as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.

point-to-point — Configures the interface to maintain this link as a point-to-point link.

message-digest-key

Syntax **message-digest-key** *keyid* **md5** [*key* | *hash-key*] [**hash**]
no message-digest-key *keyid*

Context config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-link

Description This command configures a message digest key when MD5 authentication is enabled on the interface. Multiple message digest keys can be configured.

The **no** form of the command removes the message digest key identified by the *key-id*.

Default No message digest keys are defined.

Parameters **keyid —** The *keyid* is expressed as a decimal integer.

Values 1 — 255

md5 *key* **—** The MD5 key. The *key* 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** *metric*
no metric

Context config>service>vprn>ospf>area>interface

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** *bytes*
no mtu

Context config>service>vprn>ospf>area>interface

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>interface

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>interface

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>interface
config>service>vprn>ospf>area>virtual-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

transit-delay

Syntax **transit-delay** *seconds*
no transit-delay

Context config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-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.

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

Default **transit-delay 1**

Parameters *seconds* — The transit delay in seconds expressed as a decimal integer.

Values 0 — 3600

nssa

Syntax [**no**] **nssa**

Context config>service>vprn>ospf>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

Description This command enables the generation of a default route and its LSA type (3 or 7) into a Not So Stubby Area (NSSA) by an NSSA Area Border Router (ABR)

When configuring an NSSA with no summaries, the ABR will inject a type 3 LSA default route into the NSSA area. Some older implementations expect a type 7 LSA default route.

The **no** form of the command disables origination of a default route.

Default **no originate-default-route** — A default route is not originated.

Parameters **type-7 —** Specifies a type 7 LSA should be used for the default route.

Configure this parameter to inject a type-7 LSA default route instead the type 3 LSA into the NSSA configured with no summaries.

To revert to a type 3 LSA, enter **originate-default-route** without the **type-7** parameter.

Default **Type 3 LSA for the default route.**

redistribute-external

Syntax [**no**] **redistribute-external**

Context config>service>vprn>ospf>area>nssa

Description This command enables the redistribution of external routes into the Not So Stubby Area (NSSA) or an NSSA area border router (ABR) that is exporting the routes into non-NSSA areas.

NSSA or Not So Stubby Areas are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is that the NSSA has the capability to flood external routes that it learns (providing it is an ASBR) throughout its area and via an Area Border Router to the entire OSPF domain.

The **no** form of the command disables the default behavior to automatically redistribute external routes into the NSSA area from the NSSA ABR.

Default **redistribute-external** — External routes are redistributed into the NSSA.

summaries

Syntax [**no**] **summaries**

Context config>service>vprn>ospf>area>nssa
config>service>vprn>ospf>area>stub

Description This command enables sending summary (type 3) advertisements into a stub area or Not So Stubby Area (NSSA) on an Area Border Router (ABR).

This parameter is particularly useful to reduce the size of the routing and Link State Database (LSDB) tables within the stub or nssa area.

By default, summary route advertisements are sent into the stub area or NSSA.

The **no** form of the command disables sending summary route advertisements and, for stub areas, only the default route is advertised by the ABR.

Default **summaries** — Summary routes are advertised by the ABR into the stub area or NSSA.

stub

Syntax [**no**] **stub**

Context config>service>vprn>ospf>area

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

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

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 1) 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).

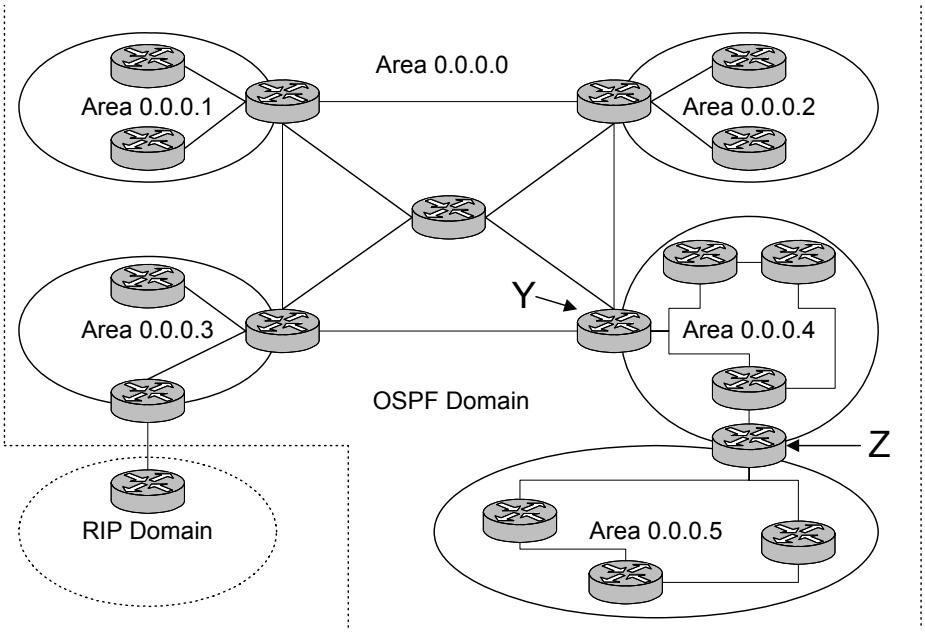

Figure 1: OSPF Areas

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.

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

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

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 *limit* — The maximum number of non-default AS-external-LSA entries that can be stored in the LSDB before going into an overflow state expressed as a decimal integer.

Values 0 — 2147483674

interval — The number of seconds after entering an overflow state before attempting to process non-default AS-external-LSAs expressed as a decimal integer.

Values 0 — 2147483674

external-preference

Syntax **external-preference** *preference*
no external-preference

Context config>service>vprn>ospf

Description This command configures the preference for OSPF external routes.

A route can be learned by the router from different protocols in which case the costs are not comparable; when this occurs the preference is used to decide which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the Table 6, “Route Preference Defaults by Route Type,” on page 1095. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the config>router context.

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

Default **external-preference 150** — OSPF external routes have a default preference of 150.

Parameters *preference* — The preference for external routes expressed as a decimal integer. Defaults for different route types are listed in Table 6.

Table 6: Route Preference Defaults by Route Type

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

overload

Syntax **overload** [**timeout** *seconds*]
no overload

Context config>service>vprn>ospf

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

overload-include-stub

Syntax [**no**] **overload-include-stub**

Context config>service>vprn>ospf

Description This command is used to to determine if the OSPF stub networks should be advertised with a maximum metric value when the system goes into overload state for any reason. When enabled, the system uses the maximum metric value. When this command is enabled and the router is in overload, all stub interfaces, including loopback and system interfaces, will be advertised at the maximum metric.

Default **no overload-include-stub**

overload-on-boot

Syntax **overload-on-boot** [**timeout** *seconds*]
no overload

Context config>service>vprn>ospf

Description When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occur:

- The timeout timer expires.
- A manual override of the current overload state is entered with the **no overload** command.

  The **no overload** command does not affect the **overload-on-boot** function.

The **no** form of the command removes the overload-on-boot functionality from the configuration.

Default **no overload-on-boot**

Parameters **timeout** *seconds* — Specifies the number of seconds to reset overloading.

Values 60 —1800

Default 60

preference

Syntax **preference** *preference*
no preference

Context config>service>vprn>ospf

This command configures the preference for OSPF internal routes.

A route can be learned by the router from different protocols in which case the costs are not comparable, when this occurs the preference is used to decide to which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in Table 7. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the config>router context.

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

Default **preference 10** — OSPF internal routes have a preference of 10.

Parameters *preference* — The preference for internal routes expressed as a decimal integer. Defaults for different route types are listed in Table 7.

Table 7: Route Preference Defaults by Route Type

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** *reference-bandwidth*
no reference-bandwidth

Context config>service>vprn>ospf

Description This command configures the reference bandwidth in kilobits per second (Kbps) that provides the reference for the default costing of interfaces based on their underlying link speed.

The default interface cost is calculated as follows:

cost = reference-bandwidth ÷ bandwidth

The default *reference-bandwidth* is 100,000,000 Kbps or 100 Gbps, so the default auto-cost metrics for various link speeds are as as follows:

- 10 Mbs link default cost of 10000
- 100 Mbs link default cost of 1000
- 1 Gbps link default cost of 100
- 10 Gbps link default cost of 10

The **reference-bandwidth** command assigns a default cost to the interface based on the interface speed. To override this default cost on a particular interface, use the **metric** *metric* command in the config>router>ospf>area>interface *ip-int-name* context.

The **no** form of the command reverts the reference-bandwidth to the default value.

Default **reference-bandwidth 100000000** — Reference bandwidth of 100 Gbps.

Parameters *reference-bandwidth* — The reference bandwidth in kilobits per second expressed as a decimal integer.

Values 1 — 1000000000

timers

Syntax **timers**

Context config>service>vprn>ospf

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

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 1 — 120000

Default 1000

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

lsa-arrival

Syntax **lsa-arrival** *lsa-arrival-time*
no lsa-arrival

Context config>service>vprn>ospf

Description This parameter defines the minimum delay that must pass between receipt of the same Link State Advertisements (LSAs) arriving from neighbors.

It is recommended that the neighbors configured (**lsa-generate**) *lsa-second-wait* interval is equal or greater then the **lsa-arrival** timer configured here.

Use the **no** form of this command to return to the default.

Default **no lsa-arrival**

Parameters *lsa-arrival-time* — Specifies the timer in milliseconds. Values entered that do not match this requirement will be rejected.

Values 0 — 600000

lsa-generate

Syntax **lsa-generate** *max-lsa-wait* [*lsa-initial-wait* [*lsa-second-wait*]]
no lsa-generate-interval

Context config>service>vprn>ospf

Description This parameter customizes the throttling of OSPF LSA-generation. Timers that determine when to generate the first, second, and subsequent LSAs can be controlled with this command. Subsequent LSAs are generated at increasing intervals of the *lsa-second-wait* timer until a maximum value is reached.

Configuring the **lsa-arrival** interval to equal or less than the *lsa-second-wait* interval configured in the **lsa-generate** command is recommended.

Use the **no** form of this command to return to the default.

Default **no lsa-generate**

Parameters *max-lsa-wait* — Specifies the maximum interval, in milliseconds, between two consecutive ocurrences 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

RIP Commands

rip

Syntax [**no**] **rip**

Context config>service>vprn

Description This command enables the RIP protocol on the given VPRN IP interface.

The **no** form of the command disables the RIP protocol from the given VPRN IP interface.

Default **no rip**

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash | hash2**]
no authentication-key

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command sets the authentication password to be passed between RIP neighbors.

The authentication type and authentication key must match exactly for the RIP message to be considered authentic and processed.

The **no** form of the command removes the authentication password from the configuration and disables authentication.

Default **no authentication-key** — Authentication is disabled and the authentication password is empty.

Parameters *authentication-key* — The authentication key. The key can be any combination of ASCII characters up to 16 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

hash-key — The hash key. The key can be any combination of ASCII characters up to 33 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** 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** {**none** | **password** | **message-digest**}
no authentication-type

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command defines the type of authentication to be used between RIP neighbors. The type and password must match exactly for the RIP message to be considered authentic and processed.

The **no** form of the command removes the authentication type from the configuration and effectively disables authentication.

Default **no authentication-type**

Parameters *none* — No authentication is used.

simple — A simple clear-text password is sent.

md5 — MD5 authentication is used.

check-zero

Syntax **check-zero** {**enable | disable**}
no check-zero

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command enables checking for zero values in fields specified to be zero by the RIPv1 and RIPv2 specifications.

The **no** form of the command disables this check and allows the receipt of RIP messages even if the mandatory zero fields are non-zero.

Default **no check-zero**

Parameters **enable —** Enables checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications and rejecting non-compliant RIP messages.

disable — Disables the checking and allows the receipt of RIP messages even if the mandatory zero fields are non-zero.

split-horizon

Syntax **split-horizon** {**enable** | **disable**}
no split-horizon

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command enables the use of split-horizon.

RIP uses split-horizon with poison-reverse to protect from such problems as “counting to infinity”. Split-horizon with poison reverse means that routes learned from a neighbor through a given interface are advertised in updates out of the same interface but with a metric of 16 (infinity).

The **split-horizon disable** command enables split horizon without poison reverse. This allows the routes to be re-advertised on interfaces other than the interface that learned the route, with the advertised metric equaling an increment of the metric-in value.

This configuration parameter can be set at three levels: global level (applies to all groups and neighbor interfaces), group level (applies to all neighbor interfaces in the group) or neighbor level (only applies to the specified neighbor interface). The most specific value is used. In particular if no value is set (**no split-horizon**), the setting from the less specific level is inherited by the lower level.

The **no** form of the command disables split horizon command which allows the lower level to inherit the setting from an upper level.

Default **enabled**

export

Syntax **export** *policy* [*policy*...]
no export

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command specifies the export policies to be used to control routes advertised to RIP neighbors. By default, RIP advertises routes from other RIP routes but does not advertise any routes from other protocols unless directed by an export policy.

The **no** form of the command removes all route policy names from the export list.

Default **no export**

Parameters *policy* — A route policy statement name.

import

Syntax **import** *policy* [*policy*...]
no import

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command specifies the import policies to be used to control routes advertised from RIP neighbors. By default, RIP accepts all routes from configured RIP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

The **no** form of the command removes all route policy names from the import list.

Default **no import**

Parameters *policy* — A route policy statement name.

message-size

Syntax **message-size** *max-num-of-routes*
no message-size

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command sets the maximum number of routes per RIP update message.

The **no** form of the command resets the maximum number of routes back to the default of 25.

Default **no message-size**

Parameters *size* — Integer.

Default 25

Values 25 — 255

metric-in

Syntax **metric-in** *metric*
no metric-in

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command sets the metric added to routes that were received from a RIP neighbor.

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

Default **no metric-in**

Parameters *metric* — The value added to the metric of routes received from a RIP neighbor, expressed as a decimal integer.

Values 1 — 16

metric-out

Syntax **metric-out** *metric*
no metric-out

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command sets the metric added to routes that were exported into RIP and advertised to RIP neighbors.

The **no** form of the command removes the command from the config and resets the metric-in value back to the default.

Default **no metric-out**

Parameters *metric* — The value added to the metric for routes exported into RIP and advertised to RIP neighbors, expressed as a decimal integer.

Values 1 — 16

preference

Syntax **preference** *preference*
no preference

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command sets the route preference assigned to RIP routes. This value can be overridden by route policies.

The **no** form of the command resets the *preference* to the default.

Default **no preference**

Parameters *preference* — An integer.

Values 1 — 255

Default 100

receive

Syntax **receive** {**both** | **none** | **version-1** | **version-2**}
no receive

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command configures the type(s) of RIP updates that will be accepted and processed.

If **both** or **version-2** is specified, the RIP instance listens for and accepts packets sent to the broadcast and multicast (224.0.0.9) addresses.

If **version-1** is specified, the router only listens for and accepts packets sent to the broadcast address.

This control can be issued at the global, group or interface level. The default behavior accepts and processes both RIPv1 and RIPv2 messages.

The **no** form of the command resets the type of messages accepted to both.

Default **no receive** — Accepts both formats.

Parameters **both —** Receive RIP updates in either Version 1 or Version 2 format.

none — Do not accept and RIP updates.

version-1 — Router should only accept RIP updates in Version 1 format.

version-2 — Router should only accept RIP updates in Version 2 format.

send

Syntax **send** {**broadcast** | **multicast** | **none** | **version-1** | **both**}
no send

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command specifies the type of RIP messages sent to RIP neighbors.

This control can be issued at the global, group or interface level. The default behavior sends RIPv2 messages with the multicast (224.0.0.9) destination address.

If **version-1** is specified, the router only listens for and accepts packets sent to the broadcast address.

The **no** form of this command resets the type of messages sent back to the default value.

Default **no send —** Sends RIPv2 to the broadcast address.

Parameters **broadcast —** Send RIPv2 formatted messages to the broadcast address.

multicast — Send RIPv2 formatted messages to the multicast address.

none — Do not send any RIP messages (i.e. silent listener).

verstion-1 — Send RIPv1 formatted messages to the broadcast address.

both — Send both RIP v1 & RIP v2 updates to the broadcast address.

timers

Syntax **timers** *update timeout flush*
no timers

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command sets the values for the update, timeout, and flush timers.

- Update timer — Determines how often RIP updates are sent.
- Timeout timer — If a router is not updated by the time the timer expires, the route is declared invalid, but maintained in the RIP database.
- Flush timer — Determines how long a route is maintained in the RIP database, after it has been declared invalid. Once this timer expires it is flushed from the RIP database completely.

The **no** form of the command resets all timers to their default values of 30, 180, and 120 seconds respectively.

Default **no timers**

Parameters *update* — The RIP update timer value in seconds.

Values 1 — 600

Default 30

timeout — The RIP timeout timer value in seconds.

Values 1 — 1200

Default 180

flush — The RIP flush timer value in seconds.

Values 1 — 1200

Default 120

group

Syntax [**no**] **group** *group-name*

Context config>service>vprn>rip

Description This command creates a context for configuring a RIP group of neighbors.

RIP groups are a way of logically associating RIP neighbor interfaces to facilitate a common configuration for RIP interfaces.

The **no** form of the command deletes the RIP neighbor interface group. Deleting the group will also remove the RIP configuration of all the neighbor interfaces currently assigned to this group.

Default **no group** — No group of RIP neighbor interfaces defined

Parameters *group-name* — The RIP group name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

neighbor

Syntax [**no**] **neighbor** *ip-int-name*

Context config>service>vprn>rip>group

Description This command creates a context for configuring a RIP neighbor interface.

By default, interfaces are not activated in any interior gateway protocol such as RIP unless explicitly configured.

The **no** form of the command deletes the RIP interface configuration for this interface. The **shutdown** command in the config>router>rip>group *group-name*>neighbor *ip-int-name* context can be used to disable an interface without removing the configuration for the interface.

Default **no neighbor** — No RIP interfaces 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.

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.

Table 8: 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.

Label	Description
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** [**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 Displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

Parameters **ingress —** Specifies matching an ingress policy.

egress — Specifies matching an egress policy.

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

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 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

Values

qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

bundle-id — Specifies the multilink bundle to be associated with this IP interface. The **bundle** keyword must be entered at the beginning of the parameter.
The command syntax must be configured as follows:

bundle-id: **bundle**-*type-slot-id*/*mda-slot*.*bundle-num*
bundle-id value range: 1 — 128

For example:

```
*A:ALA-12>config#  port bundle-ima-5/1.1
*A:ALA-12>config>port# multilink-bundle
```

ima — Specifies Inverse Multiplexing over ATM. An IMA group is a collection of physical links bundled together and assigned to an ATM port.

qtag1, qtag2 — Specifies the encapsulation value used to identify the SAP on the port or sub-port. If this parameter is not specificially defined, the default value is 0.

Values qtag1: 0 — 4094
qtag2 : * | 0 — 4094

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 TDM	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 TDM	BCP-Dot1q	0 — 4094	The SAP is identified by the 802.1Q tag on the channel.
SONET/SDH TDM	Frame Relay	16 — 991	The SAP is identified by the data link connection identifier (DLCI).
SONET/SDH ATM	ATM	vpi (NNI) 0 — 4095 vpi (UNI) 0 — 255 vci 1, 2, 5 — 65535	The SAP is identified by the PVC identifier (vpi/vci).

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.

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.
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
-------------------------------------------------------------------------------
5/1/1:0/11 511111 2 none 2 none none Up Up
5/1/1:0/12 511112 2 none 2 none none Up Up
5/1/1:0/13 511113 2 none 2 none none Up Up
5/1/1:0/14 511114 2 none 2 none none Up Up
5/1/1:0/15 511115 2 none 2 none none Up Up
5/1/1:0/16 511116 2 none 2 none none Up Up
5/1/1:0/17 511117 2 none 2 none none Up Up
5/1/1:0/18 511118 2 none 2 none none Up Up
5/1/1:0/19 511119 2 none 2 none none Up Up
5/1/1:0/20 511120 2 none 2 none none Up Up
5/1/1:0/21 511121 2 none 2 none none Up Up
5/1/1:0/22 511122 2 none 2 none none Up Up
5/1/1:0/23 511123 2 none 2 none none Up Up
5/1/1:0/24 511124 2 none 2 none none Up Up
5/1/1:0/25 511125 2 none 2 none none Up Up ...
===============================================================================
*A:ALA-12#
```

sdp

Syntax **sdp** [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]

Context show>service

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 desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Opr MTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
IP address	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Adm Admin State	Specifies the desired 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.

Label	Description
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified.
Hello Time	Specifies how often the SDP Echo Request messages are transmitted on this SDP.
Deliver Delivered	Specifies the type of delivery used by the SDP: GRE or MPLS.
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified.
Hello Time	Specifies how often the SDP Echo Request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP Echo Request messages transmitted on this SDP.
Hello Timeout	Specifies the number of seconds to wait for an SDP echo response message before declaring a timeout.
Unmatched Replies	Specifies the number of SDP unmatched message replies.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
TX Hello Msgs	Specifies the number of SDP echo request messages transmitted since the keepalive was administratively enabled or the counter was cleared.
Rx Hello Msgs	Specifies the number of SDP echo request messages received since the keepalive was administratively enabled or the counter was cleared.
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
===============================================================================
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:ALA-12# show service sdp 8
===============================================================================
Service Destination Point (Sdp Id : 8)
===============================================================================
SdpId    Adm MTU   Opr MTU   IP address       Adm  Opr         Deliver Signal
-------------------------------------------------------------------------------
8        4462      4462      10.10.10.104     Up   Dn NotReady MPLS    TLDP
===============================================================================
*A:ALA-12#
===============================================================================
Service Destination Point (Sdp Id : 8) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 8  -(10.10.10.104)
-------------------------------------------------------------------------------
Description           : MPLS-10.10.10.104
SDP Id               : 8
Admin Path MTU       : 0                    Oper Path MTU      : 0
Far End              : 10.10.10.104         Delivery           : MPLS
Admin State          : Up                   Oper State         : Down
Flags                : SignalingSessDown TransportTunnDown
Signaling            : TLDP                 VLAN VC Etype      : 0x8100
Last Status Change   : 02/01/2007 09:11:39  Adv. MTU Over.     : No
Last Mgmt Change     : 02/01/2007 09:11:46
```

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

Associated LSP LIST :
Lsp Name             : to-104
Admin State          : Up                    Oper State         : Down
Time Since Last Tran*: 01d07h36m
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:ALA-12#
```

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.
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 [epipe] [ies] [vpls] [vprn] [mirror] [apipe] [fpipe] [ipipe]** [**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 **epipe —** Displays matching Epipe services.

ies — Displays matching IES instances.

vpls — Displays matching VPLS instances.

vprn — Displays matching VPRN services.

mirror — Displays mirror services.

apipe — Displays matching Apipe services.

fpipe — Displays matching Fpipe services.

ipipe — Displays matching Ipipe 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#

de14# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
1            uVPLS     Up     Up         1                 10/26/2006 15:44:57
2            Epipe     Up     Down       1                 10/26/2006 15:44:57
10           mVPLS     Down   Down       1                 10/26/2006 15:44:57
11           mVPLS     Down   Down       1                 10/26/2006 15:44:57
100          mVPLS     Up     Up         1                 10/26/2006 15:44:57
101          mVPLS     Up     Up         1                 10/26/2006 15:44:57
102          mVPLS     Up     Up         1                 10/26/2006 15:44:57
999          uVPLS     Down   Down       1                 10/26/2006 16:14:33
-------------------------------------------------------------------------------
Matching Services : 8
-------------------------------------------------------------------------------
de14#
```

id

Syntax **id** *service-id* {**all** | **arp** | **base** | **fdb** | **labels** | **mfib** | **sap** | **sdp** | **split-horizon-group** | **stp**}

Context show>service

Description Display 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.

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.
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 multiplath 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.

Label	Description (Continued)
SAP Count	The number of SAPs specified for this service.
SDP Bind Count	The number of SDPs bound to this service.
Split Horizon Group specifics	
Split Horizon Group	Name of the split horizon group for this service.
Description	Description of the split horizon group.
Last Changed	The date and time of the most recent management-initiated change to this split horizon group.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.
Type	Indicates whether this Service SDP binding is a spoke or a mesh.
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the keepalive protocol.
Oper State	The current status of the keepalive protocol.
Hello Time	Specifies how often the SDP Echo Request messages are transmitted on this SDP.

Label	Description (Continued)
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field. 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 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 of.
Ingress sched-policy	Indicates the ingress QoS scheduler for the SAP.

Label	Description (Continued)
Egress sched-policy	Indicates the egress QoS scheduler for the SAP.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
SAP Statistics	
Dropped	The number of packets or octets dropped.
Offered Hi Priority	The number of high priority packets, as determined by the SAP ingress QoS policy.
Offered Low Priority	The number of low priority packets, as determined by the SAP ingress QoS policy.
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
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.

Label	Description (Continued)
Egress Queue 1	The index of the egress QoS queue of the SAP.
In profile for-warded	The number of in-profile packets or octets (rate below CIR) for-warded.
In profile dropped	The number of in-profile packets or octets dropped for the SAP.
Out profile for-warded	The number of out-of-profile packets or octets (rate above CIR) for-warded.
Out profile dropped	The number of out-of-profile packets or octets discarded.
DHCP Relay	
State	Specifies whether DHCP Relay is enabled on this SAP
Info Option	Specifies whether Option 82 processing is enabled on this SAP
Action	Specifies the Option 82 processing on this SAP or interface: keep, replace or drop
Circuit ID	Specifies whether the If Index is inserted in Circuit ID suboption of Option 82
Remote ID	Specifies whether the far-end MAC address is inserted in Remote ID suboption of Option 82
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.

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: 11/28/2006 12:31:53
```

```
Last Mgmt Change  : 11/28/2006 12:31:56
Admin State       : Down                Oper State        : Down

Route Dist.       : 10001:1
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         : 1                   SDP Bind Count    : 0
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------
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              : ServiceAdminDown
                     PortOperDown
Last Status Change : 11/28/2006 12:31:53
Last Mgmt Change   : 11/28/2006 12:31:56
Admin MTU          : 1518                    Oper MTU          : 1518
Ingress qos-policy : 1                       Egress qos-policy : 1
Ingress Filter-Id  : n/a                     Egress Filter-Id  : n/a

Multi Svc Site     : None
Acct. Pol          : None                    Collect Stats     : Disabled

Anti Spoofing      : None                    Nbr Static Hosts  : 0
-------------------------------------------------------------------------------
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 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
-------------------------------------------------------------------------------
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-ce1
Admin State       : Up                  Oper State        : Down
Protocols         : None

IP Addr/mask      : 11.1.0.1/24         Address Type      : Primary
IGP Inhibit       : Disabled            Broadcast Address : Host-ones
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Description  :
If Index          : 2                   Virt. If Index    : 2
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

Proxy ARP Details
Rem Proxy ARP     : Disabled            Local Proxy ARP   : Disabled
Policies          : none

DHCP Details
Admin State       : Down                Lease Populate    : 0
Action            : Keep                Trusted           : Disabled

Subscriber Authentication Details
```

```
Auth Policy       : None

ICMP Details
Redirects    : Number - 100                    Time (seconds)   - 10
Unreachables : Number - 100                    Time (seconds)   - 10
TTL Expired  : Number - 100                    Time (seconds)   - 10
===============================================================================
A:ALA-48#
```

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 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.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port*.*channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda*.*bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]

	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

Sample Output

```
*A:ALA-1# show service id 11 authentication statistics
===============================================================
Authentication statistics
===============================================================
Interface / SAP                Authentication  Authentication
                               Successful      Failed
---------------------------------------------------------------
vpls-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.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1.vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200

qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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#
```

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	The type of service: Epipe, Ipipe, Fpipe, Apipe, VPLS, IES, VPRN.
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.

Label	Description
Type	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
AdmMTU	Specifies the 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:ALA-12# show service id 1 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPRN
Customer Id       : 1
Last Status Change: 02/01/2007 09:11:39
Last Mgmt Change  : 02/01/2007 09:11:46
Admin State       : Up                  Oper State        : Down

Route Dist.       : 10001:1
AS Number         : 10000               Router Id         : 10.10.10.103
ECMP              : Enabled             ECMP Max Routes   : 8
Max Routes        : No Limit            Auto Bind         : LDP
Vrf Target        : target:10001:1
Vrf Import        : vrfImpPolCust1
Vrf Export        : vrfExpPolCust1
SAP Count         : 1                   SDP Bind Count    : 18

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                      Type         AdmMTU  OprMTU  Adm     Opr
-------------------------------------------------------------------------------
sap:1/1/7:0                     q-tag        1518    1518    Up      Up
sdp:10:1 M(10.20.1.3)           TLDP         4462    4462    Up      TLDP Down
sdp:20:1 M(10.20.1.4)           TLDP         4462    4462    Up      TLDP Down
sdp:30:1 M(10.20.1.5)           TLDP         4462    4462    Up      TLDP Down
sdp:40:1 M(10.20.1.20)          TLDP         1534    4462    Up      Up
sdp:200:1 M(10.20.1.30)         TLDP         1514    4462    Up      Up
sdp:300:1 M(10.20.1.31)         TLDP         4462    4462    Up      TLDP Down
sdp:500:1 M(10.20.1.50)         TLDP         4462    4462    Up      TLDP Down
===============================================================================
*A:ALA-12#
```

dhcp

Syntax **dhcp**

Context show>service>id

Description This command enables the context to display DHCP information for the specified service.

lease-state

Syntax **lease-state** [[**sap** *sap-id*] [**sdp** [*sdp-id*[:*vc-id*]]] | [**interface** *interface-name*] | [**ip-address** *ip-address[/mask*]] | [**mac** *ieee-address*] | [**wholesaler** *service-id*]] [**detail**]

Context show>service>id>dhcp

Description This command displays DHCP lease state related information. Refer to the following for various show command output:

- Lease State Sample Output on page 1137
- Routed CO Sample Output on page 1138
- Wholesaler/Retainer Sample Output on page 1139

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200

qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp-id — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Values 1 — 4294967295

interface *interface-name* — Displays information for the specified IP interface.

ip-address *ip-address* — Displays information associated with the specified IP address.

detail — Displays detailed information.

wholesaler *service-id* **—** The VPRN service ID of the wholesaler. When specified in this context, SAP, SDP, interface, IP address and MAC parameters are ignored.

Values 1 — 2147483647

Sample Output

Sample Output

```
*A:ALA-48>config# show service id 101 dhcp lease-state
===============================================================================
DHCP lease state table, service 101
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
102.1.1.52      00:00:1f:bd:00:bb lag-1:101           00h02m56s  DHCP-R
103.3.2.62      00:00:1f:bd:00:c6 lag-1:105           00h02m59s  Radius
-------------------------------------------------------------------------------
Number of lease states : 2
===============================================================================
*A:ALA-48>config#


*A:ALA-48>config# show service id 105 dhcp lease-state  wholesaler 101
===============================================================================
DHCP lease state table, service 105
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Wholesaler 101 Leases
-------------------------------------------------------------------------------
103.3.2.62      00:00:1f:bd:00:c6 lag-1:105           00h00m39s  Radius
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
```

```
*A:ALA-48>config#
```

Routed CO Sample Output

```
A:ALA-_Dut-A# show service id 13 dhcp lease-state
===============================================================================
DHCP lease state table, service 13
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
13.13.40.1      00:00:00:00:00:13 1/1/1:13            00h00m58s  Radius
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-_Dut-A#


A:ALA-_Dut-A# show service id 13 dhcp lease-state detail
===============================================================================
DHCP lease states for service 13
===============================================================================
Service ID           : 13
IP Address           : 13.13.40.1
Mac Address          : 00:00:00:00:00:13
Subscriber-interface : ies-13-13.13.1.1
Group-interface      : intf-13
SAP                  : 1/1/1:13
Remaining Lifetime   : 00h00m58s
Persistence Key      : N/A

Sub-Ident            : "Belgacom"
Sub-Profile-String   : "ADSL GO"
SLA-Profile-String   : "BE-Video"
Lease ANCP-String    : ""

Sub-Ident origin     : Radius
Strings origin       : Radius
Lease Info origin    : Radius

Ip-Netmask           : 255.255.0.0
Broadcast-Ip-Addr    : 13.13.255.255
Default-Router       : N/A
Primary-Dns          : 13.13.254.254
Secondary-Dns        : 13.13.254.253

ServerLeaseStart     : 12/24/2006 23:48:23
ServerLastRenew      : 12/24/2006 23:48:23
ServerLeaseEnd       : 12/24/2006 23:49:23
Session-Timeout      : 0d 00:01:00
DHCP Server Addr     : N/A

Persistent Relay Agent Information
  Circuit Id         : ancstb6_Dut-A|13|intf-13|0|13
  Remote Id          : stringtest
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-_Dut-A#
```

Wholesaler/Retainer Sample Output

```
A:ALA-_Dut-A# show service id 2000 dhcp lease-state detail
===============================================================================
DHCP lease states for service 2000
===============================================================================
-------------------------------------------------------------------------------
Wholesaler 1000 Leases
-------------------------------------------------------------------------------
Service ID           : 1000
IP Address           : 13.13.1.254
Mac Address          : 00:00:00:00:00:13
Subscriber-interface : whole-sub
Group-interface      : intf-13
Retailer             : 2000
Retailer If          : retail-sub
SAP                  : 1/1/1:13
Remaining Lifetime   : 00h09m59s
Persistence Key      : N/A

Sub-Ident            : "Belgacom"
Sub-Profile-String   : "ADSL GO"
SLA-Profile-String   : "BE-Video"
Lease ANCP-String    : ""

Sub-Ident origin     : Retail DHCP
Strings origin       : Retail DHCP
Lease Info origin    : Retail DHCP

Ip-Netmask           : 255.255.0.0
Broadcast-Ip-Addr    : 13.13.255.255
Default-Router       : N/A
Primary-Dns          : N/A
Secondary-Dns        : N/A

ServerLeaseStart     : 12/25/2006 00:29:41
ServerLastRenew      : 12/25/2006 00:29:41
ServerLeaseEnd       : 12/25/2006 00:39:41
Session-Timeout      : 0d 00:10:00
DHCP Server Addr     : 10.232.237.2

Persistent Relay Agent Information
  Circuit Id         : 1/1/1:13
  Remote Id          : stringtest
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:ALA-_Dut-A#
```

statistics

Syntax **statistics** [**sap** *sap-id*]
statistics [**sdp** *sdp-id*:*vc-id*]
statistics [**interface** *interface-name*]

Context show>service>id>dhcp

Description Displays DHCP statistics information.

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

Values *sap-id*: null [*port-id* | *bundle-id* | *lag-id* | *aps-id*]
sap-id: null[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q [*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq [*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm [*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]
frame [*port-id* | *bundle-id*]:*dlci*
cisco-hdlc *slot/mda/port.channel*

port-id *slot/mda/port*[.*channel*]
aps-id aps-*group-id*[.*channel*]
aps keyword
group-id 1 — 64
bundle-*type-slot/mda.bundle-num*
bundle keyword
type ima, ppp
bundle-num 1 — 128
bpgrp-id: **bpgrp**-*type-bpgrp-num*
bpgrp keyword
type ima
bpgrp-num 1 — 1280
ccag-id ccag-*id.path-id*[*cc-type*]:*cc-id*
ccag keyword
id 1 — 8
path-id a, b
cc-type .sap-net, .net-sap]
cc-id 0 — 4094
lag-id lag-*id*
lag keyword
id 1 — 200

qtag1 0 — 4094
qtag2 *, 0 — 4094
vpi NNI 0 — 4095
UNI 0 — 255
vci 1, 2, 5 — 65535
dlci 16 — 1022

sdp-id — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Values 1 — 4294967295

interface *interface-name* — Displays information for the specified IP interface.

Show DHCP Statistics Output — The following table describes the output fields for DHCP statistics.

Label	Description
Received Packets	The number of packets received from the DHCP clients.
Transmitted Packets	The number of packets transmitted to the DHCP clients.
Received Malformed Packets	The number of corrupted/invalid packets received from the DHCP clients.
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients. In this case, a frame is dropped due to the client sending a DHCP packet with Option 82 filled in before “trust” is set under the DHCP interface command.
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded.
Client Packets Relayed	The number of packets received from the DHCP clients that were forwarded.
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped.
Server Packets Discarded	The number of packets received from the DHCP server that were discarded.
Server Packets Relayed	The number of packets received from the DHCP server that were forwarded.
Server Packets Snooped	The number of packets received from the DHCP server that were snooped.

Sample Output

```
A:sim1# show service id 11 dhcp statistics

===============================================================
DHCP Global Statistics, service 11
===============================================================
Rx Packets                        : 32
Tx Packets                        : 12
Rx Malformed Packets              : 0
Rx Untrusted Packets              : 0
Client Packets Discarded          : 0
Client Packets Relayed            : 11
Client Packets Snooped            : 21
Server Packets Discarded          : 0
Server Packets Relayed            : 0
Server Packets Snooped            : 0
===============================================================
A:sim1#
```

gsmp

Syntax **gsmp**

Context show>service>id

Description This command displays 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                          Neigbor           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                          Neigbor           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                          Neigbor           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), neigbor 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), neigbor 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
===============================================================================
ANCP-String                                                  Assoc.
State
-------------------------------------------------------------------------------
7330-ISAM-E47 atm 1/1/01/01:19425.64048                      ANCP    Up
-------------------------------------------------------------------------------
Number of ANCP-Strings : 1
===============================================================================
A:active>show>service>id>gsmp#
```

host

Syntax **host** [**sap** *sap-id*] [**detail**]
host summary
host [**detail**] **wholesaler** *service-id*

Context show>service>id

Description This command displays static host information configured on this service.

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

summary — Displays summary host information.

wholesaler *service-id* **—** The service ID of the wholesaler.

Values 1 — 2147483647

summary

Syntax **summary**

Context show>service>id>dhcp

Description Displays DHCP configuration summary information.

Show Service-ID DHCP Summary — The following table describes show service-id DHCP summary output fields:

Label	Description
Sap/Sdp	The configuration identification, expressed by a string containing “card/mda/port/:logical-id”.
Snoop	Yes — The packets received from the DHCP clients were snooped.
	No - The packets received from the DHCP clients were not snooped.
Used/Provided	Used — The number of lease-states that are currently in use on a specific interface, that is, the number of clients on that interface got an IP address by DHCP. This value is always less than or equal to the ‘Provided’ field.
	Provided — The lease-populate value that is configured for a specific interface.
Arp Reply Agent	Displays whether or not there is proper handling of received ARP requests from subscribers.
Info Option	Keep — The existing information is kept on the packet and the router does not add any additional information..
	Replace — On ingress, the existing information-option is replaced with the information-option from the router.
	Drop — The packet is dropped and an error is logged.
Admin State	Indicates the administrative state.

Sample Output

```
A:ALA-49# show service id 1 dhcp summary
===============================================================================
DHCP Summary, service 1
===============================================================================
Interface Name                   Arp      Used/                Info    Admin
  SapId/Sdp                      Populate Provided             Option  State
-------------------------------------------------------------------------------
SpokeSDP                         No       0/0                  Keep    Down
  sdp:spoke-3:4                           0/0
test                             No       0/0                  Keep    Down
  sap:9/1/4:50/5                         0/0
to-ce1                           No       0/0                  Keep    Up
  sap:1/1/10:1                            0/0
-------------------------------------------------------------------------------
Interfaces: 3
===============================================================================
A:ALA-49#
```

interface

Syntax **interface** [*ip-address* | *ip-int-name*] [**detail**]

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.

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.
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.

Label	Description
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  :
If Index     : 26                             Virt. If Index   : 26
SAP Id       : 7/1/1.2.2
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       : 10/1/2:0
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#
```

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#
```

wholesalers

Syntax **wholesalers**

Context show>service>id

Description This command displays service wholesaler information.

Sample Output

```
*A:ALA-48>config# show service id 102 wholesalers
===============================================================================
Wholesalers for service 102
===============================================================================
Wholesaler Svc ID             Num Static Hosts         Num Dynamic Hosts
-------------------------------------------------------------------------------
101                           3                        1
-------------------------------------------------------------------------------
Number of wholesalers : 1
===============================================================================
*A:ALA-48>config#
```

Wholesaler information can also be displayed in the lease-state context.

```
*A:ALA-48>config# show service id 105 dhcp lease-state  wholesaler 101
===============================================================================
DHCP lease state table, service 105
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Wholesaler 101 Leasesok

-------------------------------------------------------------------------------
103.3.2.62      00:00:1f:bd:00:c6 lag-1:105           00h00m39s  Radius
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:ALA-48>config#
```

sap

Syntax **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 form *slot*/*mda*/*port*[**.***channel*].

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type*-*slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type*-*bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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 SAP and qtag.
Encap	The encapsulation type of the SAP.
Ethertype	Specifies an Ethernet type II Ethertype value.
Admin State	The administrative state of the SAP.
Oper State	The operating state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP.
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.
Forwarding Engine Stats	
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.
Off. Uncolor	The number of uncolored packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Queueing Stats (Ingress QoS Policy)	

Label	Description
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.
Queuing Stats (Egress QoS Policy)	
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.

Sample Output

```
*A:ALA-12# show service id 321 sap 2/1/4:0
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 321
SAP                : 2/1/4:0                Encap             : q-tag
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100

Admin State        : Up                      Oper State        : Down
Flags              : PortOperDown
                     SapIngressQoSMismatch
Last Status Change : 02/03/2007 12:58:37
Last Mgmt Change   : 02/03/2007 12:59:10
Admin MTU          : 1518                    Oper MTU          : 1518
Ingress qos-policy : 100                     Egress qos-policy : 1
Ingress Filter-Id  : n/a                     Egress Filter-Id  : n/a
```

```
Multi Svc Site     : None
Acct. Pol          : None                       Collect Stats     : Disabled
===============================================================================
*A:ALA-12#


*A:ALA-12# show service id 321 sap 2/1/4:0 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 321
SAP                : 2/1/4:0                  Encap             : q-tag
Dot1Q Ethertype    : 0x8100                    QinQ Ethertype    : 0x8100

Admin State        : Up                        Oper State        : Down
Flags              : PortOperDown
                     SapIngressQoSMismatch
Last Status Change : 02/03/2007 12:58:37
Last Mgmt Change   : 02/03/2007 12:59:10
Admin MTU          : 1518                      Oper MTU          : 1518
Ingress qos-policy : 100                       Egress qos-policy : 1
Ingress Filter-Id  : n/a                       Egress Filter-Id  : n/a
Multi Svc Site     : None
Acct. Pol          : None                       Collect Stats     : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                       Packets                  Octets
Forwarding Engine Stats
Dropped              : 0                        0
Off. HiPrio          : 0                        0
Off. LowPrio         : 0                        0
Off. Uncolor         : 0                        0

Queueing Stats(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
,,,
-------------------------------------------------------------------------------
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 Hori-zon Group	Name of the split horizon group that the SDP belongs to.
VC Type	Displays the VC type: ether, vlan, or vpls.
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.)

Label	Description (Continued)
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: GRE or MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the keepalive process.
Oper State	The operational state of the keepalive process.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
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.
E. Fwd. Octets	Specifies the number of forwarded egress octets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. If the SDP type is GRE, then the following message displays: SDP delivery mechanism is not MPLS

Sample Output

```
*A:ALA-12#  show service id 9000 sdp 2:22 detail
===============================================================================
Service Destination Point (Sdp Id : 2:22) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 2:22  -(10.10.10.103)
-------------------------------------------------------------------------------
Description     : GRE-10.10.10.103
SDP Id            : 2:22                      Type              : Spoke
Split Horiz Grp : (DSL-group1
VC Type           : Ether                     VC Tag            : n/a
Admin Path MTU    : 4462                      Oper Path MTU     : 4462
Far End           : 10.10.10.103              Delivery          : GRE
Admin State       : Up                        Oper State        : TLDP Down
Ingress Label     : 0                         Egress Label      : 0
Ingress Filter    : n/a                       Egress Filter     : n/a
Last Changed      : 10/29/2006 11:48:20       Signaling         : TLDP

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
E. Fwd. Pkts.     : 0                         E. Fwd. Octets    : 0

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS


-------------------------------------------------------------------------------
Rstp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move          : Disabled
Rstp Admin State  : Up                        Rstp Oper State   : Down
Core Connectivity : Down
Port Role         : N/A                       Port State        : Discarding
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
Designated Bridge : N/A                       Designated Port Id: 0
Active Protocol   : N/A

Fwd Transitions   : 0                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd    : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd    : 0                         TCN BPDUs tx      : 0
RST BPDUs rcvd    : 0                         RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALA-12#
```

subscriber-hosts

Syntax **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-address*[/*mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**detail**]
subscriber-hosts [**detail**] **wholesaler** *service-id*

Context show>service>id

Description This command displays subscriber host information.

Parameters **sap** *sap-id* — Displays the specified subscriber host SAP information.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

ip-address/mask — 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.

Values Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /30 subnets).
mask: 1 — 32

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

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

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

detail — Displays detailed information.

wholesaler *service-id* **—** The VPRN service ID of the wholesaler.

Values 1 — 2147483647

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#
```

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.
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.

Label	Description
Reuse time	The time when the route can be reused.
Path	The BGP AS path for the route.
Applied Policy	The applied route policy name.

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
                                                      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
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:

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.

Label	Description (Continued)
State	Disabled — The BGP peer group has been operationally disabled.
	Down — The BGP peer group is operationally inactive.
	Up — The BGP peer group is operationally active.
Peer AS	The configured or inherited peer AS for the specified peer group.
Local AS	The configured or inherited local AS for the specified peer group.
Local Address	The configured or inherited local address for originating peering for the specified peer group.
Loop Detect	The configured or inherited loop detect setting for the specified peer group.
Connect Retry	The configured or inherited connect retry timer value.
Authentication	None — No authentication is configured.
	MD5 — MD5 authentication is configured.
Local Pref	The configured or inherited local preference value.
MED Out	The configured or inherited MED value assigned to advertised routes without a MED attribute.
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.
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.

Label	Description (Continued)
Aggregator ID 0	Disabled — BGP is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group.
	Enabled — BGP is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group.
Remove Private	Disabled — BGP will not remove all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group.
	Enabled — BGP removes all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group.
Damping	Disabled — The peer group is configured not to dampen route flaps
	Enabled — The peer group is configured to dampen route flaps
Export Policy	The configured export policies for the peer group.
Import Policy	The configured import policies for the peer group.
Hold Time	The configured hold time setting.
Keep Alive	The configured keepalive setting.
Cluster Id	None — No cluster ID has been configured
	The configured route reflector cluster ID.
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.

Sample Output

```
*A:ALA-12# show router 3 bgp group
===============================================================================
BGP Groups
===============================================================================
-------------------------------------------------------------------------------
Group            : To_AS_40000
-------------------------------------------------------------------------------
Description      : Not Available
```

```
Group Type       : No Type             State           : Up
Peer AS          : 40000               Local AS        : 65206
Local Address    : n/a                 Loop Detect     : Ignore
Export Policy    : direct2bgp
Hold Time        : 90                  Keep Alive      : 30
Cluster Id       : None                Client Reflect  : Enabled
NLRI             : Unicast             Preference      : 170

List of Peers
- 10.0.0.1       : To_Jukebox
- 10.0.0.12      : Not Available
- 10.0.0.13      : Not Available
- 10.0.0.14      : To_ALA-1
- 10.0.0.15      : To_H-215

Total Peers      : 5                   Established     : 2
===============================================================================
*A:ALA-12#
```

neighbor

Syntax **neighbor** [*ip-address* [[**family** *family*] *filter1*]]
neighbor [*as-number* [[**family** *family*] filter2]]

Context show>router>bgp

Description This command displays BGP neighbor information. This command can be entered with or without any parameters.

When this command is issued without any parameters, information about all BGP peers displays.

When the command is issued with a specific IP address or ASN, information regarding only that specific peer or peers with the same AS display.

When either **received-routes** or **advertised-routes** is specified, then the routes received from or sent to the specified peer is listed (see second output example).
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 additional 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.
State	Idle — The BGP peer is not accepting connections.
	Active — BGP is listening for and accepting TCP connections from this peer.
	Connect — BGP is attempting to establish a TCP connection from this peer.
	Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer.
	Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION.
	Established — BGP has successfully established a peering and is exchanging routing information.
Last State	Idle — The BGP peer is not accepting connections.
	Active — BGP is listening for and accepting TCP connections from this peer.
	Connect — BGP is attempting to establish a TCP connection with this peer.
	Connect — BGP is attempting to establish a TCP connections from this peer.
	Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer.
	Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION.

Label	Description
Last Event	start – BGP has initialized the BGP neighbor.
	stop – BGP has disabled the BGP neighbor.
	open – BGP transport connection opened.
	close – BGP transport connection closed.
	openFail – BGP transport connection failed to open.
	error – BGP transport connection error.
	connectRetry – Connect retry timer expired.
	holdTime – Hold time timer expired.
	keepAlive – Keepalive timer expired.
	recvOpen – Receive an OPEN message.
	revKeepalive – Receive an KEEPALIVE message.
	recvUpdate – Receive an UPDATE message.
	recvNotify – Receive an NOTIFICATION message.
	None – No events have occurred.
Last Error	Displays the last BGP error and subcode 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.
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.

Label	Description
Authentication	None — No authentication is configured.
	MD5 — MD5 authentication is configured.
Next Hop Self	Disabled — BGP is not configured to send only its own IP address as the BGP nexthop in route updates to the specified neighbor.
	Enabled — BGP will send only its own IP address as the BGP nexthop in route updates to the neighbor.
AggregatorID Zero	Disabled — The BGP Neighbor is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates.
	Enabled — The BGP Neighbor is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates.
Remove Private	Disabled — BGP will not remove all private AS numbers from the AS path attribute, in updates sent to the specified neighbor.
	Enabled — BGP will remove all private AS numbers from the AS path attribute, in updates sent to the specified neighbor.
Passive	Disabled — BGP will actively attempt to establish a BGP connection with the specified neighbor.
	Enabled — BGP will not actively attempt to establish a BGP connection with the specified neighbor.
Prefix Limit	No Limit — No route limit assigned to the BGP peer group.
	1 - 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	None — No cluster ID has been configured
	The configured route reflector cluster ID.
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.

Label	Description
Recd. Prefixes	The number of routes received from the BGP neighbor.
Active Prefixes	The number of routes received from the BGP neighbor and active in the forwarding table.
Recd. Paths	The number of unique sets of path attributes received from the BGP neighbor.
Suppressed Paths	The number of unique sets of path attributes received from the BGP neighbor and suppressed due to route damping.
Input Queue	The number of BGP messages to be processed.
Output Queue	The number of BGP messages to be transmitted.
i/p Messages	Total number of packets received from the BGP neighbor.
o/p Messages	Total number of packets sent to the BGP neighbor.
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 State       : Connect
Last Event       : openFail
Last Error       : Hold Timer Expire
Connect Retry    : 20                  Local Pref.      : 100
Min Route Advt.  : 30                  Min AS Orig.     : 15
Multipath        : 1                   Multihop         : 5
Damping          : Disabled            Loop Detect      : Ignore
MED Out          : No MED Out          Authentication   : None
Next Hop Self    : Disabled            AggregatorID Zero: Disabled
Remove Private   : Disabled            Passive          : Disabled
Prefix Limit     : No Limit
Hold Time        : 90                  Keep Alive       : 30
Active Hold Time : 0                   Active Keep Alive: 0
Cluster Id       : None                Client Reflect   : Enabled
Preference       : 170                 Num of Flaps     : 0
Recd. Prefixes   : 0                   Active Prefixes  : 0
Recd. Paths      : 0                   Suppressed Paths : 0
Input Queue      : 0                   Output Queue     : 0
i/p Messages     : 0                   o/p Messages     : 0
i/p Octets       : 0                   o/p Octets       : 0
i/p Updates      : 0                   o/p Updates      : 0
Export Policy    : direct2bgp

===============================================================================
*A:ALA-12#
```

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.

Label	Description (Continued)
Flag	u – used
	s – suppressed
	h – history
	d – decayed
	* – valid
	i – igp
	e – egp
	? – incomplete
	> – best
Network	Route IP prefix and mask length for the route.
Next Hop	BGP nexthop for the route.
LocalPref	BGP local preference path attribute for the route.
MED	BGP Multi-Exit Discriminator (MED) path attribute for the route.
AS Path	The BGP AS path for the route.

Sample Output

```
*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
===============================================================================
-------------------------------------------------------------------------------
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-ex1*]
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*]

rd	{ip-address:number1 as-number1:number2 as-number2:number3}
number1	1 — 65535
as-number1	1 — 65535
number2	0 — 4294967295
as-number2	1 — 4294967295
number3	0 — 65535
ip-address	a.b.c.d
mask	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:x (eight 16-bit pieces) x:x:x:x:x:x:d.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.
Nexthop	The BGP nexthop.
From	The advertising BGP neighbor's IP address.
Res. Nexthop	The resolved nexthop.
Local Pref.	The local preference value.
Flag	u – used
	s – suppressed
	h – history
	d – decayed
	* – valid
	i – igp
	e – egp
	? – incomplete
	> – best
Aggregator AS	none – No aggregator AS attributes are present.
	The aggregator AS value.
Aggregator	none – no Aggregator attributes are present.
	The aggregator attribute value.

Label	Description
Atomic Aggr.	Atomic — The atomic aggregator flag is set.
	Not Atomic — The atomic aggregator flag is not set.
MED	none — No MED metric is present.
	The MED metric value.
Community	The BGP community attribute list.
Cluster	The route reflector cluster list.
Originator Id	none — The originator ID attribute is not present.
	The originator ID path attribute value.
Peer Router Id	The router ID of the advertising router.
AS-Path	The BGP AS path attribute.
VPRN Imported	Displays the VPRNs where a particular BGP-VPN received route has been imported and installed.

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 100.0.0.0/30 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        : 100.0.0.0/30
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                  Interface Name: to-sr7
Aggregator AS  : none                 Aggregator    : none
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#
```

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:

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured local AS setting, if not configured it is the same as the system AS.
BGP Admin State	Down — BGP is administratively disabled.
	Up — BGP is administratively enabled.
BGP Oper State	Down — BGP is operationally disabled.
	Up — BGP is operationally enabled.
Confederation AS	The configured confederation AS.
Member Confeder-ations	The configured members of the BGP confederation.
Number of Peer Groups	The total number of configured BGP peer groups.
Number of Peers	The total number of configured BGP peers.
Total BGP Active Routes	The total number of BGP routes used in the forwarding table.
Total BGP Routes	The total number of BGP routes learned from BGP peers.

Label	Description
Total BGP Paths	The total number of unique sets of BGP path attributes learned from BGP peers.
Total Path Memory	Total amount of memory used to store the path attributes.
Total Suppressed Routes	Total number of suppressed routes due to route damping.
Total History Routes	Total number of routes with history due to route damping.
Total Decayed Routes	Total number of decayed routes due to route damping.
Neighbor	BGP neighbor address.
AS (Neighbor)	BGP neighbor autonomous system number.
PktRcvd	Total number of packets received from the BGP neighbor.
PktSent	Total number of packets sent to the BGP neighbor.
InQ	The number of BGP messages to be processed.
OutQ	The number of BGP messages to be transmitted.
Up/Down	The amount of time that the BGP neighbor has either been established or not established depending on its current state.
State|Recv/Actv/Sent	The BGP neighbor's current state (if not established) or the number of received routes, active routes and sent routes (if established).

Sample Output

```
*A:ALA-12# show router 3 bgp summary
===============================================================================
 BGP Router ID : 10.0.0.14         AS : 65206   Local AS : 65206
===============================================================================
 BGP Admin State         : Up            BGP Oper State       : Up
 Confederation AS        : 40000
 Member Confederations   : 65205 65206 65207 65208

 Number of Peer Groups   : 2             Number of Peers      : 7
 Total BGP Active Routes : 86689         Total BGP Routes     : 116999
 Total BGP Paths         : 35860         Total Path Memory    : 2749476
 Total Supressed Routes  : 0             Total History Routes : 0
 Total Decayed Routes    : 0
===============================================================================
BGP Summary
===============================================================================
Neighbor           AS PktRcvd PktSent InQ OutQ   Up/Down State|Recv/Actv/Sent
-------------------------------------------------------------------------------
10.0.0.1        65206       5   21849   0    0 00h01m29s 32/0/86683
10.0.0.12       65206       0       0   0    0 00h01m29s Active
10.0.0.13       65206       5   10545   0   50 00h01m29s 6/0/86683
10.0.0.15       65205       0       0   0    0 00h01m29s Active
```

```
10.0.0.16       65206       5    9636   0   50 00h01m29s 6/0/86683
10.0.27.1           2       0       0   0    0 00h01m29s Active
10.0.28.1       60203   22512      15   0    0 00h01m29s 116955/86689/9
===============================================================================
*A:ALA-12#
```

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**]

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.
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.

Label	Description (Continued)
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.
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.

Label	Description (Continued)
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 (Time (seconds)). 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
===============================================================================
SR4#
```

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#
```

database

Syntax **database** [*ip-prefix* [/*mask*] [**longer**] [**peer** *ip-address*]

Context show>router>rip

Description Displays all routes in the RIP database.

Output **Show RIP Database Output —** The following table describes the output fields for the RIP route database.

Label	Description
Destination	The RIP destination for the route.
Peer	The router ID of the peer router.
NextHop	The IP address of the next hop.
Metric	The hop count to rate the value of different hops.
Tag	The value to distinguish between internal routes (learned by RIP) and external routes (learned from other protocols).
TTL	Displays how many seconds the specific route will remain in the routing table. When an entry reaches 0, it is removed from the routing table.

Label	Description
Valid	No — The route is not valid.
	Yes — The route is valid.

Sample Output

```
*A:ALA-1#  show rip database
===============================================================================
RIP Route Database
===============================================================================
Destination        Peer             NextHop         Metric  Tag     TTL   Valid
-------------------------------------------------------------------------------
180.0.0.10/32      180.1.7.15       0.0.0.0         2       0x0000  163   No
180.0.0.10/32      180.1.8.14       0.0.0.0         2       0x0000  179   No
180.0.0.14/32      180.1.8.14       0.0.0.0         1       0x0000  179   Yes
180.0.6.0/24       180.1.7.15       0.0.0.0         11      0x2002  163   No
180.0.6.0/24       180.1.8.14       0.0.0.0         11      0x2002  179   No
180.0.7.0/24       180.1.7.15       0.0.0.0         11      0x2002  163   No
180.0.7.0/24       180.1.8.14       0.0.0.0         11      0x2002  179   No
180.1.5.0/24       180.1.7.15       0.0.0.0         2       0x0000  151   Yes
180.1.5.0/24       180.1.8.14       0.0.0.0         1       0x0000  167   No
180.100.17.16/30   180.1.7.15       0.0.0.0         2       0x0000  151   No
180.100.17.16/30   180.1.8.14       0.0.0.0         2       0x0000  167   No
-------------------------------------------------------------------------------
No. of Routes: 11
===============================================================================
*A:ALA-12#
```

neighbor

Syntax **neighbor** [*ip-address* | *ip-int-name*] [**detail**] [**advertised-routes**]

Context show>router>rip

Description Displays RIP neighbor interface information.

Parameters *ip-address* | *ip-int-name* — Displays information for the specified IP interface.

Default All neighbor interfaces.

advertised-routes — Displays the routes advertised to RIP neighbors. If no neighbors are specified, then all routes advertised to all neighbors are displayed. If a specific neighbor is given then only routes advertised to the given neighbor/interface are displayed.

Default Display RIP information.

Output **Standard Show RIP Neighbor Output —** The following table describes the standard command output fields for a RIP group.

Label	Description
Neighbor	The RIP neighbor interface name.

Label	Description
Adm	Down — The RIP neighbor interface is administratively down.
	Up — The RIP neighbor interface is administratively up.
Opr	Down — The RIP neighbor interface is operationally down.
	Up — The RIP neighbor interface is operationally up.
Primary IP	The Primary IP address of the RIP neighbor interface.
Send Mode	Bcast — Specifies that RIPv2 formatted messages are sent to the broadcast address.
	Mcast — Specifies that RIPv2 formatted messages are sent to the multicast address.
	None — Specifies that no RIP messages are sent (i.e., silent listener).
	RIPv1 — Specifies that RIPv1 formatted messages are sent to the broadcast address.
Recv Mode	Both — Specifies that RIP updates in either version 1 or version 2 format will be accepted.
	None — Specifies that RIP updates will not be accepted.
	RIPv1 — Specifies that RIP updates in version 1 format only will be accepted.
	RIPv2 — Specifies that RIP updates in version 2 format only will be accepted.
Metric In	The metric added to routes received from a RIP neighbor.

Sample Output

```
*A:ALA-12#  show router 3 rip neighbor
===============================================================================
RIP Neighbors
===============================================================================
Interface                       Adm  Opr  Primary IP      Send  Recv  Metric
                                                         Mode  Mode  In
-------------------------------------------------------------------------------
router-2/1                      Up   Up   10.0.3.12       None  Both  1
router-2/2                      Up   Up   10.0.5.12       BCast Both  1
router-2/3                      Up   Up   10.0.6.12       BCast Both  1
router-2/4                      Up   Up   10.0.10.12      BCast Both  1
router-2/5                      Up   Up   10.0.9.12       BCast Both  1
router-2/6                      Up   Up   10.0.17.12      None  Both  1
router-2/7                      Up   Up   10.0.16.12      None  Both  1
===============================================================================
*A:ALA-12#
```

Detailed Show RIP Neighbor Output — The following table describes the standard command output fields for a RIP group.

Label	Description
Neighbor	The RIP neighbor name.
Description	The RIP neighbor description. No Description Available indicates no description is configured.
Primary IP	The RIP neighbor interface primary IP address.
Group	The RIP group name of the neighbor interface.
Admin State	Down – The RIP neighbor interface is administratively down.
	Up – The RIP neighbor interface is administratively up.
Oper State	Down – The RIP neighbor interface is operationally down.
	Up – The RIP neighbor interface is operationally up.
Send Mode	Bcast – Specifies that RIPv2 formatted messages are sent to the broadcast address.
	Mcast – Specifies that RIPv2 formatted messages are sent to the multicast address.
	None – Specifies that no RIP messages are sent (i.e., silent listener).
	RIPv1 – Specifies that RIPv1 formatted messages are sent to the broadcast address.
Recv Mode	Both – Specifies that RIP updates in either version 1 or version 2 format will be accepted.
	None – Specifies that RIP updates will not be accepted.
	RIPv1 – Specifies that RIP updates in version 1 format only will be accepted.
	RIPv2 – Specifies that RIP updates in version 2 format only will be accepted.
Metric In	The metric value added to routes received from a RIP neighbor.
Metric Out	The value added to routes exported into RIP and advertised to RIP neighbors.
Split Horizon	Disabled – Split horizon disabled for the neighbor.
	Enabled – Split horizon and poison reverse enabled for the neighbor.
Check Zero	Disabled – Checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications are not checked allowing receipt of RIP messages even if mandatory zero fields are non-zero for the neighbor.
	Enabled – checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications and rejecting non-compliant RIP messages is enabled for the neighbor.

Label	Description
Message Size	The maximum number of routes per RIP update message.
Preference	The preference of RIP routes from the neighbor.
Auth. Type	Specifies the authentication type.
Update Timer	The current setting of the RIP update timer value expressed in seconds.
Timeout Timer	The current RIP timeout timer value expressed in seconds.
Export Policies	The export route policy that is used to determine routes advertised to all peers.
Import Policies	The import route policy that is used to determine which routes are accepted from RIP neighbors.

Sample Output

```
*A:ALA-12#  show router 3 rip peers
===============================================================================
RIP Peers
===============================================================================
Peer IP Addr      Interface Name                    Version     Last Update
-------------------------------------------------------------------------------
10.0.5.13         router-2/2                        RIPv2       0
10.0.6.16         router-2/3                        RIPv2       2
10.0.9.14         router-2/5                        RIPv2       8
10.0.10.15        router-2/4                        RIPv2       0
-------------------------------------------------------------------------------
No. of Peers: 4
===============================================================================
*A:ALA-12#


*A:ALA-12# show router 3 rip neighbor detail
===============================================================================
RIP Neighbors (Detail)
===============================================================================
Neighbor "router-2/7"
-------------------------------------------------------------------------------
Description    : No Description Available
Primary IP     : 10.0.16.12         Group        : seven
Admin State    : Up                 Oper State   : Up
Send Mode      : None               Receive Mode : Both
Metric In      : 1                  Metric Out   : 1
Split Horizon  : Enabled            Check Zero   : Disabled
Message Size   : 25                 Preference   : 100
Auth. Type     : None               Update Timer : 3
Timeout Timer  : 6                  Flush Timer  : 6
Export Policies:
    Rip2Rip
    direct2Rip
    bgp2Rip
Import Policies:
    None
===============================================================================
*A:ALA-12#
```

Sample Output

```
*A:ALA-12# show router 3 rip neighbors interface advertised-routes
===============================================================================
RIP Advertised Routes
===============================================================================
Destination         Interface         NextHop           Metric  Tag        TTL
-------------------------------------------------------------------------------
180.0.0.2/32        180.1.8.12        0.0.0.0           10      0x2002     n/a
180.0.0.5/32        180.1.8.12        0.0.0.0           10      0x2002     n/a
180.0.0.8/32        180.1.8.12        0.0.0.0           10      0x2002     n/a
180.0.0.9/32        180.1.8.12        0.0.0.0           10      0x2002     n/a
180.0.0.10/32       180.1.8.12        0.0.0.0           10      0x2002     n/a
180.0.0.11/32       180.1.8.12        0.0.0.0           10      0x2002     n/a
180.0.0.12/32       180.1.8.12        0.0.0.0           1       0x0000     n/a
180.0.0.13/32       180.1.8.12        0.0.0.0           10      0x2002     n/a
180.0.0.14/32       180.1.8.12        0.0.0.0           16      0x0000     n/a
180.0.0.15/32       180.1.8.12        0.0.0.0           2       0x0000     n/a
180.0.0.16/32       180.1.8.12        0.0.0.0           3       0x0000     n/a
-------------------------------------------------------------------------------
No. of Advertised Routes: 11
===============================================================================
*A:ALA-12#
```

peer

Syntax **peer** [*ip-int-name*]

Context show>router>rip

Description Displays RIP peer information.

Parameters *ip-int-name* — Dispays peer information for peers on the specified IP interface.

Default Display peers for all interfaces.

Output **Show RIP Peer Output —** The following table describes the command output fields for a RIP peer:

Label	Description
Peer IP Addr	The IP address of the peer router.
Interface Name	The peer interface name.
Version	The version of RIP running on the peer.
Last Update	The number of days since the last update.
No. of Peers	The number of RIP peers.

statistics

Syntax **statistics** [*ip-addr* | *ip-int-name*]

Context show>router>rip

Description Display Interface level statistics for the RIP protocol.

If no IP address or interface name is specified, then all configured RIP interfaces are displayed.

If an IP address or interface name is specified, then only data regarding the specified RIP interface is displayed.

Parameters *ip-addr* | *ip-int-name* — Displays statistics for the specified IP interface.

Output **Show RIP Statistics Output —** The following table describes the output fields for RIP statistics.

Label	Description
Learned Routes	The number of RIP-learned routes were exported to RIP neighbors.
Timed Out Routes	The number of routes that have been timed out.
Current Memory	The amount of memory used by this RIP router instance.
Maximum Memory	The amount of memory allocated for this RIP router instance.
Interface	Displays the name of each interface configured in RIP and associated RIP statistics.
Primary IP	The interface IP address.
Update Timer	The current setting of the RIP update timer value expressed in seconds.
Timeout Timer	The current RIP timeout timer value expressed in seconds.
Flush Timer	The number of seconds after a route has been declared invalid that it is flushed from the route database.
Updates Sent	Total — The total number of RIP updates that were sent.
	Last 5 Min — The number of RIP updates that were sent in the last 5 minutes.
	Last 1 Min — The number of RIP updates that were sent in the last 1 minute.
Triggered Updates	Total — The total number of triggered updates sent. These updates are sent before the entire RIP routing table is sent.
	Last 5 Min — The number of triggered updates that were sent in the last 5 minutes.
	Last 1 Min — The number of triggered updates that were sent in the last 1 minute.
Bad Packets Received	Total — The total number of RIP updates received on this interface that were discarded as invalid.
	Last 5 Min — The number of RIP updates received on this interface that were discarded as invalid in the last 5 minutes.
	Last 1 Min — The number of RIP updates received on this interface that were discarded as invalid in the last 1 minute.

Label	Description
RIPv1 Updates Received	Total — The total number of RIPv1 updates received.
	Last 5 Min — The number of RIPv1 updates received in the last 5 minutes.
	Last 1 Min — The number of RIPv1 updates received in the last 1 minute.
RIPv1 Updates Ignored	Total — The total number of RIPv1 updates ignored.
	Last 5 Min — The number of RIPv1 updates ignored in the last 5 minutes.
	Last 1 Min — The number of RIPv1 updates ignored in the last 1 minute.
RIPv1 Bad Routes	Total — The total number of bad routes received from the peer.
	Last 5 Min — The number of bad routes received from the peer in the last 5 minutes.
	Last 1 Min — The number of bad routes received from the peer in the last minute.
RIPv1 Requests Received	Total — The total number of times the router received RIPv1 route requests from other routers.
	Last 5 Min — The number of times the router received RIPv1 route requests from other routers in the last 5 minutes.
	Last 1 Min — The number of times the router received RIPv1 route requests from other routers in the last 1 minute.
RIPv1 Requests Ignored	Total — The total number of times the router ignored RIPv1 route requests from other routers.
	Last 5 Min — The number of times the router ignored RIPv1 route requests from other routers in the last 5 minutes.
	Last 1 Min — The number of times the router ignored RIPv1 route requests from other routers in the last 1 minute.
RIPv2 Updates Received	Total — The total number of RIPv2 updates received.
	Last 5 Min — The number of RIPv2 updates received in the last 5 minutes.
	Last 1 Min — The number of RIPv2 updates received in the last minute.

Label	Description
RIPv2 Updates Ignored	Total — The total number of RIPv2 updates ignored.
	Last 5 Min — The number of RIPv2 updates ignored in the last 5 minutes.
	Last 1 Min — The number of RIPv2 updates ignored in the last minute.
RIPv2 Bad Routes	Total — The total number of bad routes received from the peer.
	Last 5 Min — The number of bad routes received from the peer in the last 5 minutes.
	Last 1 Min — The number of bad routes received from the peer in the last minute.
RIPv2 Requests Received	Total — The total number of times the router received RIPv2 route requests from other routers.
	Last 5 Min — The number of times the router received RIPv2 route requests from other routers in the last 5 minutes.
	Last 1 Min — The number of times the router received RIPv2 route requests from other routers in the last minute.
RIPv2 Requests Ignored	Total — The total number of times the router ignored RIPv2 route requests from other routers.
	Last 5 Min — The number of times the router ignored RIPv2 route requests from other routers in the last 5 minutes.
	Last 1 Min — The number of times the router ignored RIPv2 route requests from other routers in the last minute.
Authentication Errors	Total — The total number of authentication errors to secure table updates.
	Last 5 Min — The number of authentication errors to secure table updates in the last 5 minutes.
	Last 1 Min — The number of authentication errors to secure table updates in the last minute.

Sample Output

```
*A:ALA-12# show router 3 rip statistics
===============================================================================
RIP Statistics
===============================================================================
Learned Routes     : 0                  Timed Out Routes   : 0
Current Memory     : 120624             Maximum Memory     : 262144

-------------------------------------------------------------------------------
Interface "to-web"
-------------------------------------------------------------------------------
Primary IP         : 10.1.1.3           Update Timer       : 30
```

```
Timeout Timer      : 180               Flush Timer        : 120

Counter                       Total          Last 5 Min      Last 1 Min
-------------------------------------------------------------------------------
Updates Sent                  0              0               0
Triggered Updates             0              0               0
Bad Packets Received          0              0               0
RIPv1 Updates Received        0              0               0
RIPv1 Updates Ignored         0              0               0
RIPv1 Bad Routes              0              0               0
RIPv1 Requests Received       0              0               0
RIPv1 Requests Ignored        0              0               0
RIPv2 Updates Received        0              0               0
RIPv2 Updates Ignored         0              0               0
RIPv2 Bad Routes              0              0               0
RIPv2 Requests Received       0              0               0
RIPv2 Requests Ignored        0              0               0
Authentication Errors         0              0               0

===============================================================================
*A:ALA-12#
```

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.

Label	Description (Continued)
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
===============================================================================
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: 1
===============================================================================
*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*]]

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 optional *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.

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.

Label	Description (Continued)
Type	BH — The static route is a black hole route. The Nexthop for this type of route is black-hole.
	ID — The static route is an indirect route, where the nexthop for this type of route is the non-directly connected next hop.
	NH — The route is a static route with a directly connected next hop. The Nexthop for this type of route is either the next hop IP address or an egress IP interface name.
Next Hop	The next hop for the static route destination.
Protocol	The protocol through which the route was learned.
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#
```

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** 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.

Label	Description (Continued)
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#
```

statistics

Syntax **statistics** [*ip-int-name* | *ip-address*]

Context show>router>dhcp

Description Display statistics for DHCP Relay and DHCP snooping.

If no IP address or interface name is specified, then all configured interfaces are displayed.

If an IP address or interface name is specified, then only data regarding the specified interface is displayed.

Parameters *ip-int-name* | *ip-address* — Displays statistics for the specified IP interface.

Output **Show DHCP Statistics Output —** The following table describes the output fields for DHCP. statistics.

Label	Description
Received Packets	The number of packets received from the DHCP clients.

Label	Description
Transmitted Packets	The number of packets transmitted to the DHCP clients.
Received Malformed Packets	The number of malformed packets received from the DHCP clients.
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients.
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded.
Client Packets Relayed	The number of packets received from the DHCP clients that were forwarded.
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped.
Server Packets Discarded	The number of packets received from the DHCP server that were discarded.
Server Packets Relayed	The number of packets received from the DHCP server that were forwarded.
Server Packets Snooped	The number of packets received from the DHCP server that were snooped.

Sample Output

```
*A:ALA-1# show router dhcp statistics
===============================================================
DHCP Global Statistics
===============================================================
Rx Packets                        : 0
Tx Packets                        : 0
Rx Malformed Packets              : 0
Rx Untrusted Packets              : 0
Client Packets Discarded          : 0
Client Packets Relayed            : 0
Client Packets Snooped            : 0
Server Packets Discarded          : 0
Server Packets Relayed            : 0
Server Packets Snooped            : 0
===============================================================
*A:ALA-1#
```

summary

Syntax **summary**

Context show>router>dhcp

Description Display the status of the DHCP Relay and DHCP snooping functions on each interface.

Output **Show DHCP Summary Output —** The following table describes the output fields for DHCP summary.

Label	Description
Interface Name	Name of the router interface.
ARP Populate	Indicates whether or nor ARP populate is enabled.
Info Option	Indicates whether Option 82 is enabled.
Admin State	Indicates the administrative status.

Sample Output

```
A:ALA-48# show router dhcp summary
===============================================================================
Interface Name                    Arp      Used/     Info    Admin
                                  Populate Provided  Option  State
-------------------------------------------------------------------------------
ies-10-10.10.1.1                  Yes      1000/8000 Keep    Up
ies-100-100.100.1.1               No       0/0       Keep    Down
ies-11-11.11.1.1                  Yes      1000/8000 Keep    Up
ies-12-12.12.1.1                  Yes      1000/8000 Keep    Up
ies-13-13.13.1.1                  Yes      1000/8000 Keep    Up
ies-14-14.14.1.1                  Yes      1000/8000 Keep    Up
ies-15-15.15.1.1                  Yes      1000/8000 Keep    Up
ies-16-16.16.1.1                  No       0/0       Keep    Down
ies-2-10.17.1.1                   No       0/0       Keep    Down
ies-8-8.8.1.1                     Yes      1000/8000 Keep    Up
ies-9-9.9.1.1                     Yes      1000/8000 Keep    Up
-------------------------------------------------------------------------------
Interfaces: 11
===============================================================================
A:ALA-48#
```

Clear Commands

arp

Syntax **arp** {**all** | *ip-address*}
arp interface [*ip-int-name* | *ip-address*]

Context clear>router

Description This command clears all or specific ARP entries.

The scope of ARP cache entries cleared depends on the command line option(s) specified.

Parameters **all —** Clears all ARP cache entries.

ip-addr — Clears the ARP cache entry for the specified IP address.

interface *ip-int-name* **—** Clears all ARP cache entries for the IP interface with the specified name.

interface *ip-addr* **—** Clears all ARP cache entries for the specified IP interface with the specified IP address.

dhcp

Syntax **dhcp**

Context clear>router

Description This command enables the context to clear and reset DHCP entities.

statistics

Syntax **statistics** [**interface** *ip-int-name* | *ip-address*]

Context clear>router>dhcp

Description Clears DHCP statistics.

forwarding-table

Syntax **forwarding-table** [*slot-number*]

Context clear>router

Description This command clears the route table on the specified IOM with the route table.

If the slot number is not specified, the command forces the route table to be recalculated.

Parameters *slot-number* — Clears the specified IOM slot.

Default all IOMs

Values 1 - 10

interface

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.

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-prefixlmask*] [**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.

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.

database

Syntax **database**

Context clear>router>rip

Description Flush all routes in the RIP database.

statistics

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

Context clear>router>rip

Description This command clears statistics for RIP neighbors.

Parameters **neighbor** {*ip-address* | *ip-int-name*} **—** Clears the statistics for the specified RIP interface.

Default Clears statistics for all RIP interfaces.

id

Syntax **id** *service-id*

Context clear>service
clear>service>statistics

Description This command clears commands for a specific service.

Parameters *service-id* — The ID that uniquely identifies a service.

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280

ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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.

Values 1 — 4294967295

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.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1*.*qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1*.*vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp

	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200

qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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 /30 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.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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

spoke-sdp

Syntax **spoke-sdp** *sdp-id*:*vc-id* **ingress-vc-label**

Context clear>service>id

Description This command clears and resets the spoke SDP bindings for the service.

Parameters *sdp-id* — The spoke SDP ID to be reset.

Values 1 — 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.

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.

dhcp

Syntax [**no**] **dhcp**

Context debug>service>id

Description This command enables the context for DHCP debugging.

The **no** form of the command disables DHCP debugging.

detail-level

Syntax **detail-level** {**low** | **medium** | **high**}
no detail-level

Context debug>service>id>dhcp

Description This command configures the DHCP tracing detail level.

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

mode

Syntax **mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}
no mode

Context debug>service>id>dhcp

Description This command configures the DHCP tracing mode.

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

sap

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

Context debug>service>id
debug>service>id>dhcp
debug>service>stp

Description This command enables STP debugging for a specific SAP.

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

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1.vpi2*]	
frame	[*port-id* | *bundle-id*]:*dlci*	
cisco-hdlc	*slot/mda/port.channel*	
port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>id
debug>service>id>dhcp
debug>service>stp

Description This command enables STP debugging for a specific SDP.

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

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 srcs 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 srcs 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 srcs 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 srcs 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 srcs 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 srcs 0"

38406 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 225.1.1.2 in mode EXCLUD
E. Num srcs 0"

38407 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 225.1.1.3 in mode EXCLUD
E. Num srcs 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 srcs 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 srcs 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

```
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 miscellany.

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: 5
        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
        Group Record 4
        Type: CHG_TO_EXCL, AuxDataLen 0, Num Sources 0
        Mcast Addr: 225.1.1.5
        Source Address List
A:BA#


*A:BA# no debug
Trace disabled for all existing and future clients
*A:BA# show debug
debug
exit
```

Versatile Service Module

In This Chapter

This chapter provides information about configuring Versatile Service Module (VSM) parameters.

Topics in this chapter include:

VSM Overview

In many instances, it is desirable to process a stream of packets from one or more subscribers through multiple features that, for one reason or another, are mutually exclusive in the 7750 SR forwarding planes. For example, multiple subscriber sites could be bridged together through a VPLS instance while requiring in-service high speed Internet access (IES). Functionality of this type can be handled several ways:

- Multiple System Solution on page 1228
- Hybrid Service Solution on page 1228
- Single System Multiple Interface Solution on page 1229
- Full Feature Internal Service Cross Connect Solution on page 1229

For the purpose exploring each of these solutions, the VPLS and IES service interconnection scenario is examined.

Multiple System Solution

The multiple system (meaning multiple boxes) solution splits the functionality between two distinct nodes. The first node performs the VPLS bridging functions while maintaining per site QoS and accounting functions. The second node connects to the first node as a destination in the VPLS service. This connection could be configured as a SAP to SAP or a pseudo-wire spoke connection.

Hybrid Service Solution

The hybrid solution merges the two services into a single, common service. This can be accomplished for our example service interconnect by either supporting a virtual IP interface in the context of a VPLS service or providing an IP-only solution that provides for multiple SAPs on a single IES IP interface.

The hybrid solution does not provide for separate accounting and QoS for packets forwarded (or routed) between the subscriber sites and the packets routed to next-hops outside the subscriber domain.

Single System Multiple Interface Solution

The single system solution retains the same SLA enforcement and accounting capabilities as the multiple system solution but with the advantage of only requiring a single chassis. This is accomplished by defining the VPLS and IES services on different physical interfaces of the same type. Both interfaces are defined as access types and use the same encapsulation type (i.e., Dot1q). The services are configured with the same encapsulation values and the physical interfaces are interconnected using an external jumper cable. To avoid single point of failure issues, Link Aggregation Groups (LAG) can be used to provide an N-to-1 redundancy mechanism (as well as adding more interconnect bandwidth).

Full Feature Internal Service Cross Connect Solution

The internal service cross connect solution provides similar functionality as the single system multiple interface solution while attempting to minimize the cost, density and provisioning issues inherent to the external port jumper method. The internal service cross connection feature uses new service provisioning objects and a new type of hardware adapter to manage internal service cross connections. The remainder of this document describes the internal service cross connection feature.

Functional Components

The internal service cross connection feature uses a new adapter designed to fit within an IOM (Input Output Module) MDA (Media Dependant Adapter) slot. One or more adapters are placed into a cross connect aggregation group (CCAG). To cross connect two services, each service is bound to the same cross connect aggregation group using the same cross connection identifier. This section introduces each object and gives a brief explanation of its function.

Service Cross Connect Adapter (CCA)

The VSM Cross Connect Adapter (CCA) is a type of MDA for 7750 SR platforms designed to provide an egress to ingress forwarding plane interconnection. When a CCA is installed in an MDA slot, a set of virtual ports is available to the system providing the ability to extend packet processing through an extra set of egress and ingress forwarding paths that CCA interfaces.

Unlike external port connections which utilize two TX-RX paths, a CCA interconnects the egress forwarding path on the IOM directly to the ingress forwarding path. This eliminates the need for the physical port MAC, PHY, cable and other MDA-specific components producing a less costly and more reliable adapter. The complete 10G+ forwarding path is available allowing single conversations up to 10G.

Bandwidth is utilized more efficiently than with externally cabled ports. Typically, the offered load presented to each side of the cross-connect port-pair is asymmetric in nature. When physical ports are used to cross connect services, each service is egress bandwidth-limited to the link speed of the TX-RX path.

If one TX-RX path is under-utilized, egress services on the other path cannot make use of the available bandwidth. Since the CCA is forwarding all services over the same path, all the available bandwidth can be used.

The forwarding plane that the CCA interconnects maintains the complete egress and ingress features of the services it is interconnecting. This includes the ability to remap QoS, enforce policing and shaping, and provide ingress and egress accounting for each service.

Internal Service CCAG

VSM CCAs are placed in a CCAG. A CCAG provides a mechanism to aggregate multiple CCAs into a single forwarding group. The CCAG uses conversation hashing to dynamically distribute cross-connect traffic to the active CCAs in the aggregation group. In the event that an active CCA fails or is removed from the group, the conversation hashing function redistributes the traffic over the remaining active CCAs within the group.

The conversation hashing mechanism performed for a CCAG is identical to the hashing functions performed for Ethernet LAGs (Link Aggregation Groups).

Internal Service Cross Connect Identifier (CCID)

Services and IP interfaces are bound to a CCAG through a CCID (Cross Connect Identifier). When two services or a service and an IP interface are assigned the same CCID the CCAG attempts to provide a cross connection path between the objects. The CCID enables multiple pairs of cross connected services to share the same CCAG.

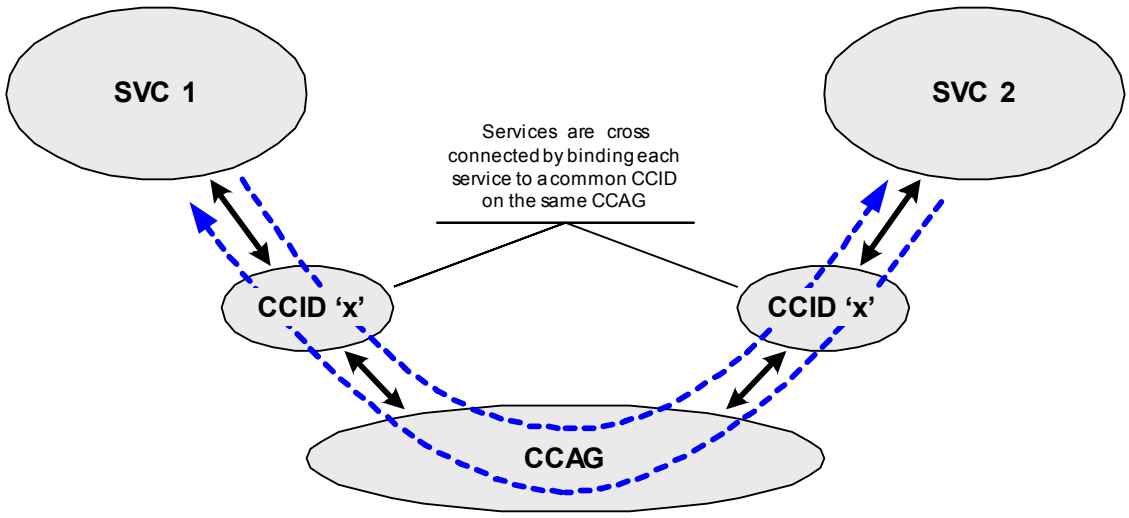

Figure 1: Internal Service Interconnection Using CCID

From a service perspective, a CCID is an object that not only binds two services together, but also provides the attachment point for the ingress and egress QoS, filtering, and accounting parameters. When considered in conjunction with the CCAG, it allows the actual cross connection path (through the CCAs) to be indirectly associated with the services using the CCAG and maintains a simplified provisioning model over port level cross connected services.

CCAG Bandwidth and Resiliency

A CCAG is an intermediate object between cross-connected objects (SAPs and network IP interfaces) and the CCAs. A CCAG is similar to a Link Aggregation Group (LAG) of Ethernet ports and uses the same underlying mechanisms to distribute conversations over multiple CCAs and converge when a CCA becomes active or inactive in the group.

When a CCAG is created, the system allocates six Ethernet LAGs for the virtual ports on the CCAs placed into the group. Each virtual port is placed into a respective LAG. For instance, each time a CCA is placed into the CCAG, virtual port 1 on that CCA is placed into the first LAG allocated to that CCAG. Virtual port 2 is placed into the second LAG on the CCAG. Virtual ports 3 through 6 are placed into their respective LAGs as well.

Using the set of LAGs provides a mechanism for conversation hashing or service mapping over all member CCAs in the CCAG. In the unlikely event that a CCA fails or is removed from the CCAG, the system will automatically modify the conversation hashing or service mapping on the CCAG to represent the available active CCAs.

CCAG LAG Attributes

Unlike user provisioned LAG, the internal LAGs do not use a primary member to control the typical port level configuration parameters. Instead, the parameters usually found at the port level are implemented directly on the CCAG internal LAG representative objects (**sap-sap**, **sap-net** and **net-sap**) for each path. These commands perform functions such as MTU definition and locally administering the MAC address.

The default unique MAC addresses used each internal LAG within the CCAG are automatically assigned from the chassis MAC pool. These MAC addresses are assigned from the pool based on an offset relative to the CCAG-ID. The same set of default MAC addresses are assigned each time a specific CCAG-ID is created.

Although a CCAG uses internal LAG mechanisms, the LACP protocol is not supported or required. LAG resources used for CCAG purposes are not exposed to the user.

CCAG Traffic Distribution

A CCAG uses both direct object mapping and conversation hashing to distribute traffic over multiple CCAs. To understand how each object type’s ingress traffic is distributed over the active CCAs in a CCAG, refer to the LAG and ECMP Hashing section of the *7750 SR OS Interface Configuration Guide*.

CCAG SAP QoS

When a SAP is created on a CCAG, the service queues defined by the ingress and egress QoS policy are created on each CCA member in the CCAG. Packets are forwarded to the egress queues based on the hashing or service mapping enforced by the LAG functions internal to the system. Packets are received on a CCA ingress queue based on which CCA handled the egress processing. Each ingress and egress hardware queue buffering and rate parameters are managed by the system based on one of two models governed by the state of the LAG QoS adaptation setting. The adaptation state also governs the application of hierarchical virtual schedulers associated with the SAP queues.

Link Level CCAG SAP QoS Adaptation

Link level QoS adaptation is set when the CCA access QoS adaptation flag is set to **link**. Link-level distribution informs the system that a service queue's buffering and rate parameters should be applied directly to each hardware queue representing the service queue. For example, when a service queue is configured with a rate equal to 10Mbps, each corresponding CCA hardware queue will be configured with a rate of 10Mbps. Given many flows conversation hashing to different CCAs, the maximum forwarded rate will be the 10Mbps multiplied by the number of active CCAs.

When a link-level adaptation service queue is a child to a parent virtual scheduler, the parent scheduler and the rest of the scheduler hierarchy is implemented per CCA. An instance of the scheduler policy is maintained per CCA.

When a CCAG SAP is a member of a Multi-Service Site (MSS), all SAPs in the MSS must be CCAG SAPs created on the same CCAG-ID.

Distributed CCAG SAP QoS Adaptation

Distributed QoS adaptation is set when the CCA access QoS adaptation flag is set to distribute. The distribute QoS parameter setting informs the system that a service queue's buffering and rate parameters should be distributed between the active CCAs in the CCAG. For example, when a service queue is configured with a rate equal to 10Mbps and two CCAs are active in the CCAG, each corresponding CCA hardware queue will be configured with a rate of 5Mbps (1/2 of the provisioned service queue parameters). Given many flows conversation hashing to different CCAs, the maximum forwarded rate will be limited to 10Mbps.

When a distributed adaptation service queue is a child to a parent virtual scheduler, the parent scheduler and the rest of the scheduler hierarchy is implemented on each IOM with an active member CCA from the CCAG. The scheduler parameters are divided amongst the IOMs with active CCAs based on the total number of active CCAs. If there are three active CCAs in the

CCAG, each CCA represents 1/3 of rate and CIR defined for each scheduler in the policy. If two of the active CCAs are on one IOM and one active CCA is on a second IOM, the first IOM would receive 2/3 of the rate and CIR for each scheduler and the second IOM would receive 1/3. The overall distribution is based on the following equation:

IOM Scheduler Rate = Policy Scheduler Rate * (Number Active CCAs on IOM / Total Active CCAs)

IOM Scheduler CIR = Policy Scheduler CIR * (Number Active CCAs on IOM / Total Active CCAs)

When a CCAG SAP is a member of a multi-service site, all SAPs in the multi-service site must be CCAG SAPs created on the same CCAG-ID.

Configuration Process Overview

Figure 58 displays the process to provision VSM parameters.

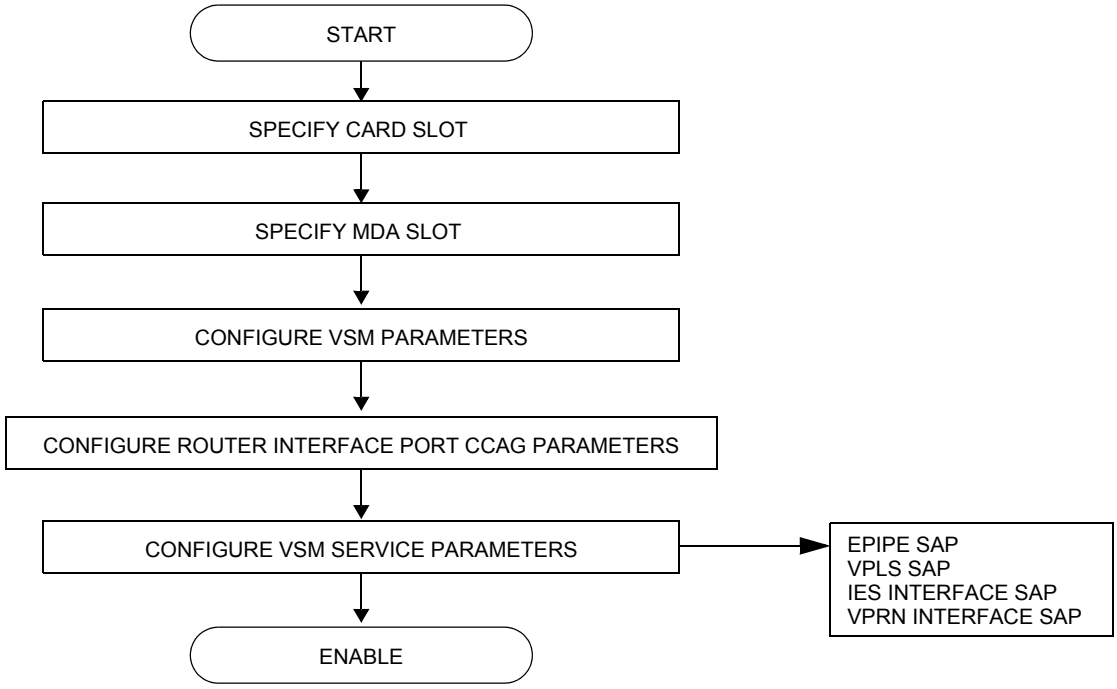

Figure 2: VSM/CCAG Configuration and Implementation Flow

Configuration Components

VSM and CCAG Components

Figure 59 displays the components to configure VSM and CCAG components.

VSM
CCAG
ACCESS
ADAPT-QOS
CCA-RATE
MEMBER-CCA
PATH
NET-SAP
RATE
SAP-NET
SAP-SAP
WEIGHT

Figure 3: VSM/CCAG Configuration Components

Figure 60 displays the basic components to configure card, router interface, and service CCAG components.

```
CONFIG
    CARD
        MDA
            MDA-TYPE
    ROUTER
        INTERFACE
            PORT
    SERVICE
        EPIPE
            SAP ccag-id
        FPIPE
            SAP ccag-id
        IPIPE
            SAP ccag-id
        VPLS
            SAP ccag-id
        IES
            INTERFACE
                SAP ccag-id
             SUBSCRIBER-INTERFACE
                 SAP ccag-id
        VPRN
            INTERFACE
                SAP ccag-id
```

Figure 4: VSM/CCAG Configuration Components

Configuration Notes

The following information describes provisioning caveats:

- Services can only be provisioned on Ethernet SAPs.
- The cross connections supported are:
  - → IP to all Layer 2 SAPs
  - → SAP to SAP of all types with the exception of:
    - – A cross connection within the same service.
    - – An IES service to another IES service.

Reference Sources

For information on supported IETF drafts and standards as well as standard and proprietary MIBS, refer to Standards and Protocol Support on page 1471.

Configuring VSM and CCAG with CLI

This section provides information to configure cards, MDAs, and ports.

Topics in this section include:

List of Commands

Table 40 lists all the configuration commands to provision VSM and service-related commands, indicating the configuration level at which each command is implemented with a short command description. The command list is organized in the following task-oriented manner:

- Configure VSM on page 1240
- Configure VSM path parameters on page 1240
- Configure VSM on an MDA on page 1241
- Configure CCAG on a router interface port on page 1241
- Configure CCAG on an Epipe service on page 1241
- Configure CCAG on a VPLS service on page 1241
- Configure CCAG on an IES service on page 1242
- Configure CCAG on an VPRN service on page 1242

Table 1: CLI Commands to Configure VSM Parameters

Command	Description	Page
Configure VSM		
`config>vsm`		
`vsm`	Changes the current CLI context to the CCA nodal context.	1272
`ccag`	Creates a Cross Connect Aggregation Group (CCAG).	1272
`access`	Changes the current CLI context to the CCAG access nodal context.	1274
`adapt-qos`	Controls how the CCAG SAP queue and virtual scheduler buffering and rate parameters are adapted over multiple active CCAs.	1274
`cca-rate`	Defines a maximum forwarding rate for each CCA member within the CCAG.	1274
`description`	Defines an informational string associated with the CCAG.	1271
`member-cca`	Adds and deletes provisioned CCAs from the CCAG.	1275
`shutdown`	Controls the administrative state of that the command is executed under.	1271
Configure VSM path parameters		
`config>vsm>path`		
`path`	Changes the current CLI context to the path nodal context.	1276
`sap-sap`	Changes the current CLI context to the path sap-sap nodal context.	1278
`sap-net`	Changes the current CLI context to the path sap-net nodal context.	1281
`net-sap`	Changes the current CLI context to the path net-sap nodal context.	1281

Table 1: CLI Commands to Configure VSM Parameters (Continued)

Command	Description	Page
mac	Overrides the default MAC address for the path's context.	1278
mtu	Overrides the default port level MTU for the path's context.	1278
egress	Changes the current CLI context to the path's egress context.	1279
ingress	Changes the current CLI context to the path's ingress context.	1280
pool	Changes the current CLI context to the path's nodal context.	1279
resv-cbs	Defines the percentage of the buffer pool that is considered reserved for the CBS buffer allocation for queues created in the path's pool context.	1279
slope-policy	Defines the slope policy used to mange the shared portion of the buffer pools WRED slopes.	1280
accounting-policy	Defines the network accounting policy that will be used to define which statistics will be collected when the **collect-stats** command is enabled in the path's net-sap context.	1282
collect-stats	Enables collecting stats on the path's net-sap context.	1282
queue-policy	Defines the egress network queues used by IP interfaces bound to the path's context	1283
rate	Defines a specific bandwidth rate limitation for the alpha or beta paths on each member CCA in the CCAG.	1276
weight	Defines a scheduling weight to the aggregate output of the alpha and beta paths.	1277
Configure VSM on an MDA		
config>card>mda		
mda	Provisions an adaptor into an MDA position on an IOM slot.	1285
Configure CCAG on a router interface port		
config>router>interface>port		
port	Cross connects a network IP interface to a CCAG SAP using the referenced *ccag-id*.	1286
Configure CCAG on an Epipe service		
config>service>epipe *service-id*		
sap	Creates a cross connect SAP on the *ccag-id* referenced in the Epipe service.	1288
Configure CCAG on a VPLS service		
config>service>vpls *service-id*		
sap	Creates a cross connect SAP on the *ccag-id* referenced in the VPLS service.	1289

Table 1: CLI Commands to Configure VSM Parameters (Continued)

Command	Description	Page
Configure CCAG on an IES service		
`config>service>ies service-id>interface ip-interface-name`		
`sap`	Creates a cross connect SAP on the *ccag-id* referenced in the IES service.	1291
Configure CCAG on an VPRN service		
`config>service>vprn service-id>interface ip-interface-name`		
`sap`	Creates a cross connect SAP on the *ccag-id* referenced in the VPRN service.	1292

Basic Configuration

The following fields require specific input (there are no defaults) to configure VSM:

- CCAG ID
- For a local service, two SAPs must be configured specifying the source and destination nodes and ports
- For a distributed service, one SAP and one SDP must be specified

The following example displays VSM defaults when a *ccag-id* is created.

```
A:ALA-48>config>vsm# info detail
#------------------------------------------
echo "Versatile Services Module Configuration"
#------------------------------------------
    vsm
        ccag 1 create
            no description
            cca-rate max
            access
                adapt-qos distribute
            exit
            path a
                weight 50
                rate max aggregate
                sap-sap
                    no mac
                    no mtu
                    egress
                        pool
                            resv-cbs default
                            slope-policy "default"
                        exit
                    exit
                    ingress
                        pool
                            resv-cbs default
                            slope-policy "default"
                        exit
                    exit
                exit
                sap-net
                    no mac
                    no mtu
                    egress
                        pool
                            resv-cbs default
                            slope-policy "default"
                        exit
                    exit
                    ingress
                        pool
                            resv-cbs default
                            slope-policy "default"
                        exit
                    exit
```

```
        exit
        net-sap
            no mac
            no mtu
            no accounting-policy
            no collect-stats
            queue-policy "default"
            egress
                pool
                    resv-cbs default
                    slope-policy "default"
                exit
            exit
        exit
    exit
    path b
        weight 50
        rate max aggregate
        sap-sap
            no mac
            no mtu
            egress
                pool
                    resv-cbs default
                    slope-policy "default"
                exit
            exit
            ingress
                pool
                    resv-cbs default
                    slope-policy "default"
                exit
            exit
        exit
        sap-net
            no mac
            no mtu
            egress
                pool
                    resv-cbs default
                    slope-policy "default"
                exit
            exit
            ingress
                pool
                    resv-cbs default
                    slope-policy "default"
                exit
            exit
        exit
        net-sap
            no mac
            no mtu
            no accounting-policy
            no collect-stats
            queue-policy "default"
            egress
                pool
                    resv-cbs default
                    slope-policy "default"
                exit
```

```
                    exit
                exit
            exit
            no shutdown
        exit
    exit
-----------------------------------------------
A:ALA-48>config>vsm#
```

Common Configuration Tasks

This section provides a brief overview of the tasks that should be performed to configure VSM on an MDA, router, router interface, and services.

- Provision one or more CCA as MDAs in the system.
- Create VSM CCAGs on the system.
- Provision CCAG path bandwidth, path weighting, and overall bandwidth parameters.
- Provision member CCAs into a CCAG.
- Provision service SAPs using a CCAG, path, and CCID for cross connect purposes.
- Bind routed network IP interfaces to a CCAG, path, and CCID for cross connect purposes.

Configure VSM CCAG Components

Use the CLI syntax displayed below to configure the following entities:

Provision VSM on an MDA

Before a CCA module may be utilized in the system, the CCA must be provisioned into an MDA slot. The MDA provisioning command must be modified to support provisioning a CCA adaptor type. Up to 8 member CCAs can be configured per CCAG.

CLI Syntax: `config>card# mda` *`mda-number mda-type`* `{other-MDA-type|cca}`

The following example displays the command usage to provision CCA on an MDA:

Example:
```
config# card 10
    config>card# mda 1
    config>card>mda# mda-type vsm-cca
    config>card>mda# exit
    config>card#
```

The following example displays the configuration:

```
A:ALA-48>config>card# info
----------------------------------------------
        card-type iom-20g
        mda 1
            mda-type vsm-cca
        exit
        mda 2
            mda-type m20-1gb-tx
        exit
----------------------------------------------
A:ALA-48>config>card#
```

Provision CCAG Parameters

Once a CCA is provisioned into the system, it must be placed in a Cross Connect Aggregation Group (CCAG) to be used by cross connect objects. Besides CCA membership, the CCAG also supports bandwidth control parameters (see Configure Path Components on page 1249) used to manipulate forwarding distribution between objects in the alpha and beta path groups and the aggregate rate allowed on the CCA.

Use the following CLI syntax to provision CCAG components.

CLI Syntax:
```
config>vsm#
            ccag ccag-id [create]
            cca-rate kilobits-per-second
            description description-string
            member-cca card-slot/mda-number
            path {a|b}
            no shutdown
```

The following example displays the command usage to provision CCAG components:

Example:
```
config>vsm# ccag 1
       config>vsm>ccag# description "VSM test"
       config>vsm>ccag# cca-rate 1000000
       config>vsm>ccag# member-cca 10/1
```

The following example displays the configuration:

```
A:ALA-48>config>vsm# info
----------------------------------------------
        ccag 1 create
            description "VSM test"
            cca-rate 1000000
            member-cca 10/1
        exit
...
----------------------------------------------
A:ALA-48>config>vsm#
```

Configure Path Components

Each CCA is divided into two distinct paths for bandwidth management purposes. One path is identified as alpha (a) and the other beta (b). The significance of each path for bandwidth distribution is dependent on the relative path weights each path is given in relationship to the other. A maximum path rate may also be defined allowing the provisioning of a maximum cap on the aggregate bandwidth allowed to the SAP or IP interface queues associated with the path.

Each path is separated into three other contexts; SAP-2-SAP (sap-sap), SAP-2-Net (sap-net) and Net-2-SAP (net-sap). Each path context allows for the definition of the features that are usually associated with physical ports on other MDAs in the system. These include buffer pool management, ingress network queue definitions and accounting policy control.

Use the following CLI syntax to provision path components.

- Net SAP
- SAP net
- SAP SAP

Use the following CLI syntax to provision CCAG path components.

CLI Syntax:

```
config>vsm>ccag#
          path {a|b}
             net-sap
                accounting-policy policy-id
                collect-stats
                egress
                   pool
                      resv-cbs percent-or-default
                      slope-policy slope-policy-name
                mac ieee-address
                mtu mtu-bytes
                queue-policy queue-policy-name
             rate kilo-bits-per-second [aggregate|cca]
             sap-net
                egress
                   pool
                      resv-cbs percent-or-default
                      slope-policy slope-policy-name
                ingress
                   pool
                      resv-cbs percent-or-default
                      slope-policy slope-policy-name
                mac ieee-address
                mtu mtu-bytes
             sap-sap
```

```
egress
   pool
      resv-cbs percent-or-default
      slope-policy slope-policy-name
ingress
   pool
      resv-cbs percent-or-default
      slope-policy slope-policy-name
mac ieee-address
mtu mtu-bytes
weight path-weight
```

The following example displays the command usage to provision CCAG path parameters:

Example:
```
config>vsm# ccag 1 create
config>vsm>ccag# path a
config>vsm>ccag>path# weight 100
config>vsm>ccag>path# exit
config>vsm>ccag# path b
config>vsm>ccag>path# rate 99999999
config>vsm>ccag>path# weight 100
config>vsm>ccag>path# exit
onfig>vsm>ccag# no shutdown
```

The following example displays the configuration:

```
A:ALA-48>config>vsm# info
----------------------------------------------
        ccag 1 create
            description "VSM test"
            member-cca 10/1
            path a
                weight 100
            exit
            path b
                weight 100
                rate 99999999
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-48>config>vsm#
```

Cross Connecting Network IP Interfaces

To support cross connection between services and network IP interfaces, the network interface port command has been augmented to allow the binding of the IP interface to a **ccag** *cc-id*. Similar to service CCAG SAPs, the network IP interface port binding command must reference the ccag-id, the CCA path (.a or .b) and the *cc-id* used by the service CCAG SAP on the other CCA path.

Use the following CLI syntax to configure CCAG a a network IP interface.

CLI Syntax:
```
config# router [router-name]
        interface interface-name
           port ccag-ccag-id.{a|b}[.net-sap]:cc-id
           address {ip-address/mask |ip-address netmask} [broadcast
              all-ones|host-ones]
           mac ieee-address
```

The following example displays the command usage:

Example:
```
config# router
config>router# interface ccanet
config>router>if$ address 2.1.1.1/24
config>router>if# port ccag-1.a.net-sap:200
config>router>if# exit
config>router# interface ccanet2
config>router>if$ address 4.1.1.1/24
config>router>if# port ccag-1.b.net-sap:300
config>router>if# static-arp 4.1.1.2 00:00:00:00:00:aa
config>router>if# exit
config>router#
```

```
A:ALA-48>config>router# info
----------------------------------------------
#----------------------------------------
echo "IP Configuration"
#----------------------------------------
...
        interface "ccanet"
            address 2.1.1.1/24
            port ccag-1.a.net-sap:200
            mac 00:00:00:00:00:ff
        exit
        interface "ccanet2"
            address 4.1.1.1/24
            port ccag-1.b.net-sap:300
            static-arp 4.1.1.2 00:00:00:00:00:aa
        exit
...
#----------------------------------------
A:ALA-48>config>router#
```

Cross Connecting Services

Services are provisioned onto a CCAG using a special CCAG SAP definition. CCAG SAPs must reference a *ccag-id*, a CCA path (a or b), a pairing type (sap-sap or sap-net) and a unique *cc-id*. The *ccag-id* identifies the group of CCAs that will be used for forwarding packets associated with the SAP. The path identifies the bandwidth control grouping used to manage CCA egress bandwidth. The pairing type helps the system identify which buffering resources will be used to manage egress queuing of packets. Finally, the *cc-id* is used to explicitly cross connect the SAP to another SAP or network IP interface configured with the same *cc-id*.

Epipe

CLI Syntax:
```
config>service#
            epipe service-id [customer customer-id]
               sap ccag-ccag-id.{a|b}[.sap-net|.sap-sap]:cc-id [create]
```

Example:
```
config>service# epipe 103 customer 6 create
       config>service>epipe$ sap 3/1/1.1.1 create
       config>service>epipe>sap# no shutdown
       config>service>epipe>sap# exit
       config>service>epipe# sap ccag-1.a:100 create
       config>service>epipe>sap$ no shutdown
       config>service>epipe>sap# exit
       config>service>epipe# no shutdown
```

The following output displays the configuration:

```
A:ALA-48>config>service# info
----------------------------------------
...
        epipe 103 customer 6 vpn 103 create
            sap 3/1/1.1.1 create
            exit
            sap ccag-1.a:100 create
            exit
            no shutdown
        exit
----------------------------------------
A:ALA-48>config>service#
```

VPLS

CLI Syntax:
```
config>service#
    vpls service-id [customer customer-id]
        sap ccag-ccag-id.{a|b}[.sap-net|.sap-sap]:cc-id [create]
```

Example:
```
config>service# vpls 740 customer 1 create
config>service>vpls# sap 1/1/19:1 create
config>service>vpls>sap$ no shutdown
config>service>vpls>sap$ exit
config>service>vpls# sap 1/1/19:2 create
config>service>vpls>sap$ no shutdown
config>service>vpls>sap>ingress$ qos 3
config>service>vpls>sap>ingress$ exit
config>service>vpls>sap# egress
config>service>vpls>sap>egress# qos 1010
config>service>vpls>sap>egress# exit
config>service>vpls>sap# exit
config>service>vpls# sap ccag-1.a:456 create
config>service>vpls>sap$ ingress
config>service>vpls>sap>ingress$ qos 3
config>service>vpls>sap>ingress$ exit
config>service>vpls>sap# egress
config>service>vpls>sap>egress# qos 1010
config>service>vpls>sap>egress# exit
config>service>vpls>sap# exit
config>service>vpls# no shutdown
```

The following output displays the configuration:

```
A:ALA-48>config>service# info
----------------------------------------
...
        vpls 740 customer 1 vpn 740 create
            stp
                shutdown
            exit
            sap 1/1/19:1 create
            exit
            sap 1/1/19:2 create
                ingress
                    qos 3
                exit
            exit
            sap ccag-1.a:456 create
                ingress
                    qos 3
                exit
                egress
                    qos 1010
                exit
            exit
            no shutdown
        exit
...
----------------------------------------
A:ALA-48>config>service#
```

IES

CLI Syntax:
```
config>service#
    ies service-id [customer customer-id]
       interface ip-interface-name
          sap ccag-ccag-id.{a|b}[.sap-net|.sap-sap]:cc-id [cre-
          ate]
```

Example:
```
config>service# ies 200 customer 1 create
config>service>ies$ interface "ccaiesif" create
config>service>ies>if$ address 8.1.1.1/24
config>service>ies>if$ sap ccag-1.b:456 create
config>service>ies>if>sap$ ingress
config>service>ies>if>sap>ingress$ qos 3
config>service>ies>if>sap>ingress$ exit
config>service>ies>if>sap# egress
config>service>ies>if>sap>egress# qos 1010
config>service>ies>if>sap>egress# exit
config>service>ies>if>sap# no shutdown
config>service>ies>if>sap# exit
config>service>ies>if# no shutdown
config>service>ies>if# exit
config>service>ies# no shutdown
config>service>ies# exit
config>service#
```

The following output displays the configuration:

```
A:ALA-48>config>service# info
----------------------------------------
...
        ies 200 customer 1 create
            interface "ccaiesif" create
                address 8.1.1.1/24
                sap ccag-1.b:456 create
                    ingress
                        qos 3
                    exit
                    egress
                        qos 1010
                    exit
                exit
            exit
            no shutdown
        exit
...
----------------------------------------
A:ALA-48>config>service#
```

VPRN

CLI Syntax:
```
config>service#
    vprn service-id [customer customer-id]
       interface ip-interface-name
          sap ccag-ccag-id.{a|b}[.sap-net|.sap-sap]:cc-id [cre-
          ate]
```

Example:
```
config>service# vprn 701 customer 2 create
config>service>vprn$ interface "VSM Test" create
config>service>vprn>if$ sap ccag-2.a:100 create
config>service>vprn>if>sap$ no shutdown
config>service>vprn>if>sap# exit
config>service>vprn>if# exit
config>service>vprn# no shutdown
```

The following output displays the configuration:

```
A:ALA-48>config>service>vprn# info
----------------------------------------------
            interface "VSM Test" create
                sap ccag-2.a:100 create
                exit
            exit
            no shutdown
----------------------------------------------
A:ALA-48>config>service>vprn#
```

Service Management Tasks

This section discusses the following service management tasks:

Modifying or Deleting a VSM MDA

To change or delete a VSM MDA already provisioned for a specific slot, first you must shut down and remove all service SAP and router interface associations (page 1257) to delete the VSM MDA from the configuration.

CLI Syntax:
```
config> card slot-number
    [no] mda mda-number
        [no] mda-type mda-type
        shutdown
```

Example:
```
config# card 10
    config>card# mda 1
    config>card>mda# mda-type vsm-cca
    config>card>mda# shutdown
    config>card>mda# exit
    config>card# no mda 1
```

The following example displays the configuration:

```
A:ALA-48>config>card# info
----------------------------------------------
        card-type iom-20g
        mda 2
            mda-type vsm-cca
        exit
----------------------------------------------
A:ALA-48>config>card#
```

Modifying CCAG Parameters on a Network IP Interface

CLI Syntax: config# router [*router-name*]
interface *interface-name*
shutdown
no port ccag-*ccag-id*.{a|b}[.net-sap]:*cc-id*

The following example displays the command usage:

Example: config>router# interface ccanet
config>router>if# address **3**.1.1.1/24
config>router>if# exit

```
A:ALA-48>config>router# info
----------------------------------------
#----------------------------------------
echo "IP Configuration"
#----------------------------------------
...
        interface "ccanet"
            address 3.1.1.1/24
            port ccag-1.a.net-sap:200
            mac 00:00:00:00:00:ff
        exit
        interface "ccanet2"
            address 4.1.1.1/24
            port ccag-1.b.net-sap:300
            static-arp 4.1.1.2 00:00:00:00:00:aa
        exit
...
#----------------------------------------
A:ALA-48>config>router#
```

Modifying CCAG Parameters

CLI Syntax: config>vsm#
ccag *ccag-id* [create]
no ccag *ccag-id* [force]
access {link|distribute}
adapt-qos
cca-rate *kilobits-per-second*
no cca-rate
description *description-string*
no description
[no] member-cca *card-slot/mda-number*
path {a|b}
no shutdown

The following example displays the command usage to provision CCAG components:

Example:
```
config>vsm# ccag 1
        config>vsm>ccag# access
        config>vsm>ccag>access#
        config>vsm>ccag>access# adapt-qos distribute
        config>vsm>ccag>access# exit
        config>vsm>ccag# member-cca 10/2
        config>vsm>ccag# exit
```

The following example displays the configuration:

```
A:ALA-48>config>vsm# info
----------------------------------------------
        ccag 1 create
            description "VSM test"
            member-cca 10/1
            member-cca 10/2
            path a
                weight 100
            exit
            path b
                weight 100
                rate 99999999
            exit
            no shutdown
        exit

...
----------------------------------------------
A:ALA-48>config>vsm# ccag 1
```

Modifying Path Parameters

CLI Syntax:

```
config>vsm>ccag#
    path {a|b}
       net-sap
          accounting-policy policy-id
          no accounting-policy
          [no] collect-stats
          egress
             pool
                resv-cbs percent-or-default
                no resv-cbs
                slope-policy slope-policy-name
                no slope-policy
          mac ieee-address
          no mac
          mtu mtu-bytes
          no mtu
          queue-policy queue-policy-name
          no queue-policy
       rate kilo-bits-per-second [aggregate|cca]
       no rate
       sap-net
          egress
             pool
                resv-cbs percent-or-default
                no resv-cbs
                slope-policy slope-policy-name
                no slope-policy
          ingress
             pool
                resv-cbs percent-or-default
                no resv-cbs
                slope-policy slope-policy-name
                no slope-policy
          mac ieee-address
          no mac
          mtu mtu-bytes
          no mtu
       sap-sap
          egress
             pool
                resv-cbs percent-or-default
                no resv-cbs
                slope-policy slope-policy-name
                no slope-policy
          ingress
             pool
                resv-cbs percent-or-default
```

```
                    no resv-cbs
                    slope-policy slope-policy-name
                    no slope-policy
                mac ieee-address
                no mac
                mtu mtu-bytes
                no mtu
            weight path-weight
            no weight
```

The following example displays the command usage to provision CCAG path parameters:

Example:
```
config>vsm# ccag 1
config>vsm>ccag# path a
config>vsm>ccag>path# no weight
config>vsm>ccag>path# net-sap
config>vsm>ccag>path>net-sap# queue-policy nq1
config>vsm>ccag>path>net-sap# egress
config>vsm>ccag>path>net-sap>egr# pool
config>vsm>ccag>path>net-sap>egr>pool# slope-policy A
config>vsm>ccag>path>net-sap>egr>pool# exit
config>vsm>ccag>path>net-sap>egr# exit
config>vsm>ccag>path>net-sap# exit
config>vsm>ccag>path# exit
config>vsm>ccag# path b
config>vsm>ccag>path# no rate
config>vsm>ccag>path# sap-sap
config>vsm>ccag>path>sap-sap# egress
config>vsm>ccag>path>sap-sap>egr# pool
config>vsm>ccag>path>sap-sap>egr>pool#
config>vsm>ccag>path>sap-sap>egr>pool# slope-policy B
config>vsm>ccag>path>sap-sap>egr>pool# exit
config>vsm>ccag>path>sap-sap>egr# exit
config>vsm>ccag>path>sap-sap# exit
config>vsm>ccag>path# exit
config>vsm>ccag#
```

The following example displays the configuration:

```
A:ALA-48>config>vsm# info
----------------------------------------------

        ccag 1 create
            description "VSM test"
            member-cca 10/1
            member-cca 10/2
            path a
                net-sap
                    queue-policy "nq1"
                    egress
```

```
                        pool
                            slope-policy "A"
                        exit
                    exit
                exit
            exit
            path b
                weight 100
                sap-sap
                    egress
                        pool
                            slope-policy "B"
                        exit
                    exit
                exit
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-48>config>vsm#
```

Modifying Service Parameters

Epipe

CLI Syntax:
```
config>service#
            epipe service-id
               sap ccag-ccag-id.{a|b}[.sap-net|.sap-sap]:cc-id
               no sap sap-id
               shutdown
```

The following service examples display the command usage to provision CCAG.

Example:
```
config>service# epipe 103
        config>service>epipe#  sap ccag-1.a:100
        config>service>epipe>sap# shutdown
        config>service>epipe>sap# exit
        config>service>epipe# no sap ccag-1.a:100
        config>service>epipe# sap ccag-1.b:200 create
        config>service>epipe>sap$ no shutdown
        config>service>epipe>sap$ exit
        config>service>epipe#
```

The following output displays the configuration:

```
A:ALA-48>config>service>epipe# info
----------------------------------------------
           sap 3/1/1.1.1 create
           exit
           sap ccag-1.b:200 create
           exit
           no shutdown
----------------------------------------------
A:ALA-48>config>service>epipe#
```

VPLS

CLI Syntax:
```
config>service#
            vpls service-id [customer customer-id]
               sap ccag-ccag-id.{a|b}[.sap-net|.sap-sap]:cc-id
               no sap sap-id
               shutdown
```

Example:
```
config>service>vpls# sap ccag-1.a:456
config>service>vpls>sap# shutdown
config>service>vpls>sap# exit
config>service>vpls# no sap ccag-1.a:456
```

```
            config>service>vpls# sap ccag-1.b:100 create
            config>service>vpls>sap$ no shutdown
            config>service>vpls>sap$ exit
            config>service>vpls# sap ccag-1.a:100
            config>service>vpls>sap# ingress
            config>service>vpls>sap>ingress# qos 3
            config>service>vpls>sap>ingress# exit
            config>service>vpls>sap# egress
            config>service>vpls>sap>egress# qos 1010
            config>service>vpls>sap>egress# exit
            config>service>vpls>sap# exit

A:ALA-48>config>service>vpls# info
----------------------------------------------
            stp
                shutdown
            exit
            sap 1/1/19:1 create
            exit
            sap 1/1/19:2 create
                ingress
                    qos 3
                exit
            exit
            sap ccag-1.b:100 create
            exit
            no shutdown
----------------------------------------------
A:ALA-48>config>service>vpls#
```

IES

CLI Syntax:

```
config>service#
   ies service-id [customer customer-id]
      interface ip-interface-name
         sap ccag-ccag-id.{a|b}[.sap-net|.sap-sap]:cc-id
         no sap sap-id
         shutdown
```

Example:

```
config>service# ies 200
config>service>ies# interface "ccanet6"
config>service>ies>if# sap ccag-1.a:101 create
config>service>ies>if>sap# ingress
config>service>ies>if>sap>ingress# qos 3
config>service>ies>if>sap>ingress# exit
config>service>ies>if>sap# egress
config>service>ies>if>sap>egress# qos 1010
config>service>ies>if>sap>egress# exit
config>service>ies>if>sap# no shutdown
config>service>ies>if>sap# exit
config>service>ies>if#
```

The following output displays the configuration:

```
A:ALA-48>config>service>ies# info
----------------------------------------------
            interface "ccaiesif" create
                address 8.1.1.1/24
                sap ccag-1.b:456 create
                    ingress
                        qos 3
                    exit
                    egress
                        qos 1010
                    exit
                exit
            exit
            interface "ccanet6" create
                address 7.1.1.1/24
                sap ccag-1.a:101 create
                    ingress
                        qos 3
                    exit
                    egress
                        qos 1010
                    exit
                exit
            exit
            no shutdown
----------------------------------------------
A:ALA-48>config>service>ies#
```

VPRN

CLI Syntax: config>service#
vprn *service-id* [customer *customer-id*]
interface *ip-interface-name*
sap ccag-*ccag-id*.{a|b}[.*sap-net*|.*sap-sap*]:*cc-id*
no sap *sap-id*
shutdown

On a VPRN service SAP:

Example:
```
config>service#   vprn 701
config>service>vprn# interface "VSM-Test Config" create
config>service>vprn>if$ sap ccag-2.b:50 create
config>service>vprn>if>sap$ no shutdown
config>service>vprn>if>sap$ exit
config>service>vprn>if# exit
config>service>vprn#
```

The following output displays the configuration:

```
A:ALA-48>config>service>vprn# info
----------------------------------------------
```

```
            interface "VSM Test" create
                sap ccag-2.a:100 create
                exit
            exit
            interface "VSM-Test Config" create
                sap ccag-2.b:50 create
                exit
            exit
            no shutdown
-----------------------------------------------
A:ALA-48>config>service>vprn#
```

VSM Command Reference

Command Hierarchies

VSM Configuration Commands

```
config
    — vsm
        — ccag ccag-id [create]
        — no ccag ccag-id [force]
            — access
                — adapt-qos {link | distribute}
                — no adapt-qos
        — cca-rate kilobits-per-second
        — no cca-rate
        — description description-string
        — no description
        — [no] member-cca card-slot/mda-number
        — path {a | b}
            — net-sap
                — accounting-policy accounting-policy
                — no accounting-policy
                — [no] collect-stats
                — egress
                    — pool
                        — resv-cbs percentage-of-pool
                        — no resv-cbs
                        — slope-policy slope-policy-name
                        — no slope-policy
                — mac mac-address
                — no mac
                — mtu mtu-size
                — no mtu
                — queue-policy queue-policy-name
                — no queue-policy
            — rate kilobits-per-second [aggregate | cca]
            — no rate
            — sap-net
                — egress
                    — pool
                        — resv-cbs percentage-of-pool
                        — no resv-cbs
                        — slope-policy slope-policy-name
                        — no slope-policy
                — ingress
                    — pool
                        — resv-cbs percentage-of-pool
                        — no resv-cbs
```

```
                                    — slope-policy slope-policy-name
                                    — no slope-policy
                            — mac mac-address
                            — no mac
                            — mtu mtu-size
                            — no mtu
                    — sap-sap
                            — egress
                                    — pool
                                    — resv-cbs percentage-of-pool
                                    — no resv-cbs
                                    — slope-policy slope-policy-name
                                    — no slope-policy
                            — ingress
                                    — pool
                                    — resv-cbs percentage-of-pool
                                    — no resv-cbs
                                    — slope-policy slope-policy-name
                                    — no slope-policy
                            — mac mac-address
                            — no mac
                            — mtu mtu-size
                            — no mtu
                    — weight path-weight
                    — no weight
            — [no] shutdown
```

Related Commands

```
config card slot-number
    — mda {1 | 2} type {existing-mda-types | vsm}
    — [no] mda {1 | 2}
config router [router-name]
    — [no] interface ip-interface-name
        — port ccag-ccag-id.{a | b}[.net-sap]:cc-id
        — no port
config service
    — epipe service-id [customer customer-id]
        — sap ccag-ccag-id.{a | b}[.sap-net | .sap-sap]:cc-id [create]
        — no sap ccag-ccag-id.{a | b}[.sap-net | .sap-sap]:cc-id
config service
    — vpls service-id [customer customer-id]
        — sap ccag-ccag-id.{a | b}[.sap-net | .sap-sap]:cc-id [create]
        — no sap ccag-ccag-id.{a | b}[.sap-net | .sap-sap]:cc-id
ies service-id [customer customer-id]
    — interface ip-interface-name
        — sap ccag-ccag-id.{a | b}[.sap-net | .sap-sap]:cc-id [create]
        — no sap ccag-ccag-id.{a | b}[.sap-net | .sap-sap]:cc-id
vprn service-id [customer customer-id]
    — interface ip-interface-name
        — sap ccag-ccag-id.{a | b}[.sap-net | .sap-sap]:cc-id [create]
```

— **no sap ccag**-*ccag-id*.{**a** | **b**}[.*sap-net* | .*sap-sap*]:*cc-id*

VSM Configuration Commands

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>vsm>ccag ccag-id

Description This command controls the administrative state of the *ccag-id* the command is executed under. Upon creation, the default state of a CCAG is to be administratively up which corresponds to the no shutdown form of the command. If the CCAG must be forced to be operationally down, the **shutdown** command will place the CCAG into an administratively down state causing the operational state to also be down.

When a CCAG is shutdown, all SAPs associated with the CCAG will be operationally down. An operationally down SAP cannot be used for forwarding packets. If the SAP is part of the VPLS service, all MAC entries associated with the SAP will be removed from the VPLS FDB and the SAP will be removed from the flooding domain of the VPLS. If the SAP is part of an IES service, the associated IP interface will be set to an operationally down state. Network IP interfaces bound to a shutdown CCAG will be operationally down as well.

Executing the no shutdown command sets the CCAG to the default up administrative state. As long as at least one member CCA in the CCAG is active, all SAPs and network IP interfaces associated with the CCAG will be allowed to enter the operationally up state.

Default **no shutdown**

description

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

Context config>vsm

Description This command defines an informational string associated with the CCAG. The description string may be up to 80 characters long and contain only printable ASCII characters. Each time this command is successfully executed, any previous description string will be overwritten. If the command fails due to improper string definition, a previously successful description string will remain.

The **no** form of the command removes any current description string from the CCAG.

Default **None** (A description string must be explicitly defined)

Parameters *description-string* — Defines the string of printable ASCII characters, up to 80 characters that will be stored and displayed as a description for the *ccag-id* that the **description** command is executed under. The string must be entered in double quotation marks if the string contains spaces.

VSM CLI Tree Node Commands

vsm

Syntax [**no**] **vsm**

Context config

Description This command changes the current CLI context to the CCA nodal context. The CCA nodal context is where CCAGs are created and maintained.

The CCA nodal context always exists and cannot be removed.

ccag

Syntax **ccag** *ccag-id* [**create**]
no ccag *ccag-id* [**force**]

Context config>vsm

Description This command creates a Cross Connect Aggregation Group (CCAG). A CCAG represents a group of CCAs as a common forwarding entity. Objects requiring a CCA cross connect function are mapped to a CCAG, not the individual CCAs within the CCAG. The CCAG treats each active member CCA as a possible destination when forwarding packets between the cross connected objects mapped to the CCAG. The system uses both conversation hashing functions and direct service mappings to determine the load sharing distribution between the active CCAs. All packets for a given conversation flow through the same CCA to preserve packet order. Packet ordering may be momentarily affected during convergence events when CCAs are dynamically added or removed from the active list.

The CCAG context is used to manage the following functions per CCAG instance:

- Informational description of the CCAG
- Administrative state of the CCAG
- Alpha path bandwidth and weight parameters
- Beta path bandwidth and weight parameters
- CCA total bandwidth limit
- CCA membership in the CCAG

The **no** form of the command removes an existing *ccag-id* from the system. Once the specified *ccag-id* is removed from the system, it may not be referenced by any cross connect objects. If the **force** keyword is not specified, the **no ccag** *ccag-id* command will fail if the specified *ccag-id* has one or more *cc-ids* associated with it. In the event that the specified *ccag-id* does not exist, the **no ccag** *ccag-id* command will return to the current CLI context without any change to the system.

Default **None** (each CCAG context must be explicitly created to be used)

Parameters *ccag-id* — Identifies the CCAG instance that the system is creating or editing. Up to eight CCAGs may be created within the system. A *ccag-id* must be created on the system prior to creating cross connect object associations.

After a *ccag-id* is created, a CCAG SAP may be created with an association with the *ccag-id*. A CCAG SAP is identified by a concatenation of an existing *ccag-id* and a *cc-id*. The *cc-id* must match the *cc-id* of the other object the CCAG SAP is paired with on the *ccag-id*. The created *ccag-id* may also be associated with a network IP interface. A network IP interface is bound to the *ccag-id* through the port command in the config router interface ip-interface context and references the *ccag-id* and a *cc-id*. Again, the *cc-id* must match the other object the IP interface is paired with on the *ccag-id*.

Once created, the **ccag** *ccag-id* command may be executed to enter the *ccag-id* instance for the purpose of editing the CCAG parameters or operational state.

Values 1 through 8

create — The **create** keyword explicitly indicates that the specified *ccag-id* is being created. Handling the inclusion or exclusion state of the create keyword is dependent on the system environment variable create.

When the system environment variable create is enabled, the system requires the explicit use of the create keyword when creating objects such as a CCAG. If the keyword is not included and the *ccag-id* has not already been created, an error will occur and the CLI will remain at the current CLI context. This is designed to prevent the inadvertent creation of a CCAG instance in the event where the wrong *ccag-id* is specified during an attempt to edit an existing CCAG instance. If the create keyword is specified, the *ccag-id* will be created given the *ccag-id* is within the proper range for CCAG identifiers.

When the system environment variable create is disabled (using the no create command), the system will not require the create keyword when creating a CCAG instance. In the event that the ccag command is issued with a *ccag-id* that previously had not been created, that *ccag-id* will be considered available for cross connect associations and bindings.

Once a *ccag-id* has been created, the create keyword is ignored when a ccag command is executed with that *ccag-id*. The **ccag** *ccag-id* create command will only result in a CLI context change to the specified CCAG instance for a pre-existing *ccag-id*.

force — The **force** keyword removes the specified *ccag-id* regardless of the presence of one or more *cc-id*. If a SAP exists on the *ccag-id*, the force keyword will cause the SAP to be removed from the configuration. If a network IP interface is bound to the *ccag-id*, the interface will be silently unbound from the *ccag-id*. The force keyword is intended as a time saving feature, preventing the need to first remove all service and network associations with the *ccag-id*.

It is not required to first remove all CCAs from the CCAG prior to deleting the CCAG from the system. When the CCAG is removed, association with all member CCAs is automatically removed.

access

Syntax **access**

Context config>vsm>ccag *ccag-id*

Description This command changes the current CLI context to the CCAG access nodal context. The access nodal context contains the qos adaptation command used to control the SAP QoS distribution across the active member CCAs within the CCAG.

The CCAG access nodal context always exists and cannot be removed.

adapt-qos

Syntax **adapt-qos** {**link** | **distribute**}
no adapt-qos

Context config>vsm>ccag *ccag-id*>access

Description This command controls how the CCAG SAP queue and virtual scheduler buffering and rate parameters are adapted over multiple active CCAs. Two adaptation modes are supported; link and distributed.

The **no** form of the command returns the CCAG access QoS adaptation rule to the default setting of distribute.

Parameters **link —** The **link** keyword is mutually exclusive with the **distributed** keyword. When **link** is specified, the CCAG will create the SAP queues and virtual schedulers on each CCA with the actual parameters defined in the QoS and scheduler policies. This mode is useful when conversation hashing places all or most traffic over a single CCA.

Values link or distribute

Default **distribute**

distribute — The **distribute** keyword is mutually exclusive with the **link** keyword. When **distribute** is specified, the CCAG SAP queues and schedulers on each CCA will receive a portion of the defined parameters in the QoS and scheduler policies. The portion is decided on an IOM basis with the ratio determined by the number of active CCA members on the IOM relative to the total number of active members within the CCAG. The following equation may be used to determine the actual ratio:

IOM-parameter-value = (IOM-active-CCA / total-active-CCA) * policy-parameter-value

Values link or distribute

Default **distribute**

cca-rate

Syntax **cca-rate** *kilobits-per-second*

no cca-rate

Context config cca>ccag *ccag-id*

Description This command defines a maximum forwarding rate for each CCA member within the CCAG. Support of setting a maximum CCA forwarding rate is provided to prevent overrunning the ingress forwarding plane when sub-line rate ingress features are enabled. The primary ingress feature requiring this support is dual ingress access queuing. When dual ingress queuing is enabled on cross connect SAPs, the CCA forwarding rate should be limited to a rate that prevents packet loss due to ingress forwarding congestion. The specified cca-rate limit is applied to the aggregate alpha and beta path bandwidth.

The **no** form of the command removes CCA bandwidth rate limiting.

Parameters *kilobits-per-second* — Defines the maximum CCA rate in kilobits per second. The actual Kilobits per second rate is rounded up to the nearest 50Mbps increment.

Values 0 — 100000000, max

Default **max**

member-cca

Syntax [**no**] **member-cca** *card-slot/mda-number*

Context config>vsm>ccag *ccag-id*

Description This command adds and deletes provisioned CCAs from the CCAG. The only requirement to defining a CCA member is that the defined MDA position be provisioned as type cca. A CCA does not need to be populated in the defined MDA position prior to membership definition. A non-populated CCA member is considered inactive from a CCAG perspective. A populated CCA member will become active once it has been initialized by the system. A CCA member may be removed from the CCAG or depopulated from MDA slot at any time. At least one member CCA must be active on the CCAG for the CCAG to be placed in the operational state. Up to 8 member CCAs can be configured per CCAG.

The **no** form of the command removes a CCA member from the CCAG. If the CCA does not exist or is not currently a member of the CCAG, no error is returned. Once removed from the CCAG, all forwarding through the specified CCA stops.

Parameters *card-slot/mda-number* — Identifies the system MDA slot that is will be added as a member CCA for the CCAG. The specified MDA slot must have been pre-provisioned as type cca for the membership command to be successful.

card-slot — Defines the IOM slot the provisioned CCA is or will be populated. It is separated from the following mda-position portion of the parameter by a forward slash (/).

Values 1 through 10 (chassis type dependent)

mda-position — The mda-position portion of the parameter defines the MDA slot number on the IOM the CCA is or will be populated. It must be separated from the preceding card-slot portion of the parameter by a forward slash (/).

Values 1 or 2 (IOM type dependent)

VSM Path Commands

path

Syntax **path** {**a** | **b**}

Context config>vsm>ccag *ccag-id*

Description This command changes the current CLI context to the path nodal context. The CCA path nodal context is where each CCA path bandwidth, buffer and accounting parameters are maintained. The path context command must be specified with either the a or b keyword specifying the CCA path context to be entered.

Each CCA is divided into two distinct paths for bandwidth management purposes. One path is identified as alpha (a) and the other beta (b). The significance of each path for bandwidth distribution is dependent on the relative path weights each path is given in relationship to the other. A maximum path rate may also be defined allowing the provisioning of a maximum cap on the aggregate bandwidth allowed to the SAP or IP interface queues associated with the path. Each path is separated into three other contexts; SAP-2-SAP (sap-sap), SAP-2-Net (sap-net) and Net-2-SAP (net-sap). Each path context allows for the definition of the features that are usually associated with physical ports on other MDAs in the system. These include buffer pool management, ingress network queue definitions and accounting policy control.

The CCA path nodal contexts always exist and cannot be removed.

Parameters **a —** The **a** keyword is mutually exclusive to the **b** keyword and defines the CLI CCA path context to be the alpha path. Either the **a** or **b** path must be specified. If the **a** or **b** keyword is not present, the path command will fail without changing the current CLI context.

b — The **b** keyword is mutually exclusive to the **a** keyword and defines the CLI CCA path context to be the beta path. Either the **a** or **b** path must be specified. If the **a** or **b** keyword is not present, the path command will fail without changing the current CLI context.

rate

Syntax **rate** *kilobits-per-second* [**aggregate** | **cca**]
no rate

Context config>vsm>ccag *ccag-id*>path {a | b}

Description This command defines a specific bandwidth rate limitation for the alpha or beta paths on each member CCA in the CCAG. Use of the rate command is optional. When the rate command is not executed or the no rate command is executed, bandwidth allocated to the path is not limited to a specific rate.

Path limiting on a CCA prevents the aggregate bandwidth for the path from exceeding a certain rate. If the rate is exceeded, the CCA will backpressure all active egress queues sending on that path. Access to the available bandwidth is dependent on the various parameters associated with each object egress queue.

The specified rate may be defined as an aggregate path rate for all CCAs in the CCAG or it may be defined as a per CCA path rate.

The **no** form of the command removes path rate limiting from all CCAs in the CCAG membership list for the path.

Default **None** (rate limiting the alpha path must be explicitly defined)

Parameters *kilobits-per-second* — Defines the path rate in kilobits per second. The **aggregate** and **cca** keywords specify how the defined rate is applied on a per CCA basis. The actual rate at each CCA is rounded up to the nearest 50Mbps.

Values 0 — 100000000, max

Default **max**

aggregate — The **aggregate** keyword is optional and mutually exclusive to the **cca** keyword. When **aggregate** is specified, the defined rate is divided among the CCAs in the CCAG member list based on the number of active CCAs. If three CCAs are active, the rate is divided by three and the result is applied to each active CCA. If a fourth CCA becomes active on the CCAG, the defined rate is then divided by four with the result applied to each CCA member on the CCAG. The actual rate at each CCA is implemented in 50Mbps increments. The system will adapt the specified rate to the best rate available per CCA.

Default When the kilobits-per-second parameter is specified, the default keyword is aggregate.

cca — The **cca** keyword is optional and mutually exclusive to the **aggregate** keyword. When **cca** is specified, the defined rate is applied to all CCAs in the CCAG member list. The actual rate at each CCA is implemented in 50Mbps increments. The system will adapt the specified rate to the best rate available per CCA.

weight

Syntax **weight** *path-weight*
no weight

Context config cca>ccag ccag-id>path {a | b}

Description This command defines a scheduling weight to the aggregate output of the alpha and beta paths. The specified weight is used to calculate a scheduling percentage for each path. The percentage for each path is based on:

Alpha scheduling percentage = alpha-path-weight / (alpha-path-weight + beta-path-weight)

Beta scheduling percentage = blue-path-weight / (alpha-path-weight + beta-path-weight)

Based on the above calculation, the sum of the alpha and beta scheduling percentage always equals 100 percent. When one path is not using all of its available scheduling bandwidth, the other path may use the remainder.

The **no** form of the command returns the path-weight for the path to the default value of 50.

Parameters *path-weight* — The path-weight parameter is required and is used by the system to determine the scheduling percentage for both paths. Changing the path-weight for one path affects both paths

scheduling percentage. The resulting scheduling percentage changes are applied to all CCAs in the CCAG membership list.

Values 1 to 100

Default **50**

sap-sap

Syntax **sap-sap**

Context config cca>ccag ccag-id>path {a | b}

Description This command changes the current CLI context to the path SAP-SAP nodal context. This context contains the ingress and egress buffer pool configuration commands. The sap-sap>path context is associated with all SAPs defined on the CCAG path (alpha or beta depending on the path context) that cross connect to a SAP on the other path.

The CCA path SAP-SAP nodal context always exists and cannot be removed.

mac

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

Context config>vsm>ccag *ccag-id*>path {a | b}
config>vsm>ccag *ccag-id*>path {a | b}>sap-net
config>vsm>ccag *ccag-id*>path {a | b}>net-sap

Description This command overrides the default MAC address for the path's context.

The **no** form of the command returns the in-use MAC address for the path's context to the default MAC from the chassis MAC pool.

Parameters *mac-address* — Defines the IEEE MAC address that is to be associated with the path's context.

Values Any valid IEEE MAC source MAC address
(6 byte address expressed in hexadecimal notation with each byte separated by a dash (-))

Default The path's default sap-sap MAC address is derived from the chassis MAC address pool

mtu

Syntax **mtu** *mtu-size*
no mtu

Context config>vsm>ccag *ccag-id*>path {a | b}>sap-sap
config>vsm>ccag *ccag-id*>path {a | b}>sap-net

config>vsm>ccag *ccag-id*>path {a | b}>net-sap

Description This command overrides the default port level MTU for the path's context.

The **no** form of the command returns the MTU for the path's sap-sap context to the default MTU.

Parameters *mtu-size* — Defines the Ethernet MTU that is to be associated with the path's context.

Default **1518** - sap-sap
1518 - sap-net
9212 - net-sap

Values 512 — 9212 bytes

egress

Syntax **egress**

Context config>vsm>ccag *ccag-id*>path {a | b}>sap-sap
config>vsm>ccag *ccag-id*>path {a | b}>sap-net
config>vsm>ccag *ccag-id*>path {a | b}>net-sap

Description This command changes the current CLI context to the path's context. This context contains the egress buffer pool configuration commands.

The CCA path's egress nodal context always exists and cannot be removed.

pool

Syntax **pool**

Context config>vsm>ccag *ccag-id*>path {a | b}>sap-sap>egress
config>vsm>ccag *ccag-id*>path {a | b}>sap-sap>ingress
config>vsm>ccag *ccag-id*>path {a | b}>sap-net>egress
config>vsm>ccag *ccag-id*>path {a | b}>sap-net>ingress
config>vsm>ccag *ccag-id*>path {a | b}>net-sap>egress

Description This command changes the current CLI context to the path's nodal context. This context contains the egress buffer pool configuration commands.

The CCA path's egress or ingress pool nodal context always exists and cannot be removed.

resv-cbs

Syntax [**no**] **resv-cbs percentage-of-pool**

Context config>vsm>ccag *ccag-id*>path {a | b}>sap-sap>egress>pool
config>vsm>ccag *ccag-id*>path {a | b}>sap-sap>ingress>pool
config>vsm>ccag *ccag-id*>path {a | b}>sap-net>egress>pool

config>vsm>ccag *ccag-id*>path {a | b}>sap-net>ingress>pool
config>vsm>ccag *ccag-id*>path {a | b}>net-sap>egress>pool

Description This command defines the percentage of the buffer pool that is considered reserved for the CBS buffer allocation for queues created in the path's pool context.

The **no** form of the command returns the reserved portion of the buffer pool to the default percentage.

Parameters *percentage-of-pool* — The percentage-of-pool parameter defines the percentage of the buffer pool that is not considered shared. The shared portion of the pool is used by queues that have crossed their CBS buffer threshold and is subject to the WRED slope functions. The reserved portion of the pool is used by queues that have not crossed their CBS threshold. The aggregate CBS on the queues associated with the pool may oversubscribe the resv-cbs percentage. If the reserved portion is oversubscribed and the in-use reserved buffers exceed the defined percentage, buffers are removed from the shared portion of the pool.

Values 1 to 100 (percent)

Default **30**

slope-policy

Syntax **slope-policy slope-policy-name**
no slope-policy

Context config>vsm>ccag *ccag-id*>path {a | b}>sap-sap>egress>pool
config>vsm>ccag *ccag-id*>path {a | b}>sap-sap>ingress>pool
config>vsm>ccag *ccag-id*>path {a | b}>sap-net>egress>pool
config>vsm>ccag *ccag-id*>path {a | b}>sap-net>ingress>pool
config>vsm>ccag *ccag-id*>path {a | b}>net-sap>egress>pool

Description This command defines the slope policy used to mange the shared portion of the buffer pools WRED slopes. The commands in the policy control the administrative state of the slopes, the start and knee points of each slope and the time-average-factor for the weighted average buffer utilization calculation.

The **no** form of the command configures the default slope policy as the managing policy for the buffer pool.

Parameters *slope-policy-name* — Defines the name of the WRED slope policy used to manage the WRED slopes in the shared portion of the buffer pool.

Values Any existing slope policy name

ingress

Syntax **ingress**

Context config>vsm>ccag *ccag-id*>path {a | b}>sap-sap
config>vsm>ccag *ccag-id*>path {a | b}>sap-net

Description This command changes the current CLI context to the path's context. This context contains the ingress buffer pool configuration commands.

The CCA path's ingress nodal context always exists and cannot be removed.

sap-net

Syntax **sap-net**

Context config>vsm>ccag *ccag-id*>path {a | b}

Description This command changes the current CLI context to the path sap-net nodal context. This context contains the ingress and egress buffer pool configuration commands. The sap-net>path context is associated with all SAPs defined on the CCAG path (alpha or beta depending on the path context) that cross connect to a network IP interface on the other path.

The CCA path sap-net nodal context always exists and cannot be removed.

slope-policy

Syntax **slope-policy** *slope-policy-name*
no slope-policy

Context config>vsm>ccag *ccag-id*>path {a | b}>sap-net>ingress>pool

Description This command defines the slope policy used to mange the shared portion of the buffer pools WRED slopes. The commands in the policy control the administrative state of the slopes, the start and knee points of each slope and the time-average-factor for the weighted average buffer utilization calculation.

The **no** form of the command configures the default slope policy as the managing policy for the buffer pool.

Parameters *slope-policy-name* — The slope-policy-name parameter defines the name of the WRED slope policy used to manage the WRED slopes in the shared portion of the buffer pool.

Values Any existing slope policy name

net-sap

Syntax **net-sap**

Context config>vsm>ccag *ccag-id*>path {a | b}>net-sap

Description This command changes the current CLI context to the path net-sap nodal context. The net-sap nodal context contains the network accounting and queue policies and the egress buffer pool configuration commands. The net-sap path context is associated with all network IP interfaces bound to the CCAG path (alpha or beta depending on the path context) that cross connects to a SAP on the other path.

The CCA path net-sap nodal context always exists and cannot be removed.

mtu

Syntax **mtu** *mtu-size*
no mtu

Context

Description This command overrides the default port level MTU for the path's net-sap context.

The **no** form of the command returns the MTU for the path's net-sap context to the default MTU.

Parameters *mtu-size* — The mtu-size, in bytes, defines the Ethernet MTU that is to be associated with the path's net-sap context.

Default **1522**

accounting-policy

Syntax **accounting-policy** *accounting-policy*
no accounting-policy

Context config>vsm>ccag *ccag-id*>path {a | b}>net-sap

Description This command defines the network accounting policy that will be used to define which statistics will be collected when the **collect-stats** command is enabled in the path's net-sap context.

The **no** form of the command reverts the path's net-sap context statistics billing collection to the statistics defined in the default network accounting policy.

Parameters *accounting-policy* — The accounting-policy parameter is required and identifies which set of statistics will be collected for billing XML output.

Values Any existing network accounting policy in the system

Default **The default network accounting policy**

collect-stats

Syntax [**no**] **collect-stats**

Context config>vsm>ccag *ccag-id*>path {a | b}>net-sap

Description This command enables collecting stats on the path's net-sap context. When enabled the statistics defined in the accounting-policy accounting-policy command will be collected according to the specifications in the policy.

The **no** form of the command disables network billing statistics collection on the net-sap context.

Default Network statistics are not collected by default on the net-sap context

queue-policy

Syntax **queue-policy** *queue-policy-name*
no queue-policy

Context config>vsm>ccag *ccag-id*>path {a | b}>net-sap

Description This command defines the egress network queues used by IP interfaces bound to the path's net-sap context. The specified *queue-policy-name* defines the number of queues, the rate and buffering parameters for the queues and the forwarding class mappings to the queues.

The **no** form of the command reverts the path's net-sap network IP interface queues to the systems default queue policy.

Parameters *queue-policy-name* — Specifies which existing Queue Policy will define the queuing structure for network IP interfaces bound to the path's net-sap context.

Values Any existing queue policy on the system

Default The default queue policy is used when another is not specified

egress

Syntax **egress**

Context config>vsm>ccag *ccag-id*>path {a | b}>net-sap

Description This command changes the current CLI context to the path>net-sap>egress nodal context. This context contains the egress buffer pool configuration commands.

The CCA path net-sap egress nodal context always exists and cannot be removed.

pool

Syntax **pool**

Context config>vsm>ccag *ccag-id*>path {a | b}>net-sap>egress

Description This command changes the current CLI context to the path>net-sap>egress pool>nodal context. This context contains the egress buffer pool configuration commands.

The CCA path net-sap egress pool nodal context always exists and cannot be removed.

Related Commands

Refer to the *7750 SR OS Interface Configuration Guide* for more card, MDA, and port command information. Refer to the *7750 SR OS Services Guide* for details about configuring specific service parameters.

mda

Syntax **mda** *mda-slot*
no mda *mda-slot*

Context config>card

Description This command provisions an adaptor into an MDA position on an IOM slot. The provisioned MDA may or may not exist in the system at the time of provisioning. If the provisioned MDA does not currently exist in the specified MDA position number, it is considered to be a 'ghost' MDA. Ports and other resources on a ghost MDA may be configured once the MDA is provisioned. When a proper MDA matching the provisioned MDA type is inserted into the IOM MDA position, forwarding though the MDA based on configured services or network interface will be available once the MDA has been properly initialized.

A Versatile Service Module (VSM) is provisioned into the system in the same manner as all other adaptors using MDA slots. Once a VSM is provisioned, independent of it actually existing in the system on the specified slot and MDA position, the VSM may be defined as a member of a CCAG (Cross Connect Adaptor Group). A VSM inserted into the system prior to provisioning is not available for CCAG membership and will be treated as an unprovisioned MDA.

Once a VSM is provisioned and populated in the system, it cannot be used until it has been defined membership into a CCAG. When the CCAG membership has been defined for the VSM, the various internal resources of the VSM will be configured according to the CCAG bandwidth control parameters. This includes the alpha and beta path weights, the alpha and beta path maximum rates and the aggregate alpha and beta maximum rate.

The **no** form of the command unprovisions an MDA from the system. For a VSM to be unprovisioned, the VSM must not be a member of a CCAG. If the VSM is a member of a CCAG, the **no cca** *slot-number/mda-number* command must be used in the CCAG member-list context. Once a CCA is unprovisioned from the system; it cannot be made a member of a CCAG until it has been reprovisioned.

Default **None** (An MDA position number must be explicitly specified.)

Parameters *mda-slot* — Defines the position on the card slot-number the CCA will be populated into. On the iom-20g IOM module, two MDA positions are available. Future IOMs may support a different number of MDA positions.

Values 1 or 2

port

Syntax **port ccag**-*ccag-id*.{**a** | **b**}[.*net-sap*]:*cc-id*
no port

Context config>router>interface *ip-interface-name*

Description This command cross connects a network IP interface to a CCAG SAP using the referenced *ccag-id*. A CCAG network IP interface binding is identified by four items; the *ccag-id*, the CCAG path, the pairing type and the *cc-id*. A network IP interface CCAG port binding supports all the available features as port binding using a Dot1Q virtual interface.

To support cross connection between services and network IP interfaces, the network interface port command allows the binding of the IP interface to a *ccag cc-id*. Similar to service CCAG SAPs, the network IP interface port binding command must reference the *ccag-id*, the CCA path (.a or .b) and the *cc-id* used by the service CCAG SAP on the other CCA path. The pairing type is optional as only *.net-sap* is supported.

The **no** form of the command removes the CCAG binding from the network IP interface.

Parameters **ccag —** The ccag portion of the port binding is required and specifies that the network IP interface is binding to a **ccag** *cc-id*.

ccag-id — The *ccag-id* portion of the port binding is required and specifies which *ccag-id* the network IP interface must be bound to. The specified *ccag-id* must exist on the system or the port binding will fail. The leading dash must be included as a separator between ccag and the *ccag-id*.

Values -1 (dash 1) to -8 (dash 8)

Default **None**

.a | .b — The .a and .b portion of the port binding is required and is used to define the CCA bandwidth path the network IP interface will be associated with. The path association must be specified and .a and .b are mutually exclusive. The .a designation identifies the network IP interface as being on the Alpha path and the .b designation identifies the network IP interface as being on the Beta path. The paired SAP using the same *cc-id* as the bound network IP interface must be associated with the opposite path. The leading period must be included as a separator between the *ccag-id* and the path designator.

Values .a or .b

Default **None**

.net-sap — The .net-sap portion of the network IP interface CCAG binding is optional and is used to explicitly define the pairing type as Net-2-SAP. A cross connection between two network IP interfaces is not currently allowed. The .net-sap pairing type is assumed and does not need to be included in the SAP identification. When specified, the leading period must be used as a separator between the path designator and the pairing type.

Default **.net-sap**

:cc-id — The *:cc-id* portion of the port binding is required and specifies the unique *cc-id* in use by the CCAG network IP interface port binding and the cross connect SAP on the other path.

Values 1 to 4094

Service CCAG SAP Provisioning

Services are provisioned onto a CCAG using a special CCAG SAP definition. CCAG SAPs must reference a *ccag-id*, a CCA path (a or b), a pairing type (sap-sap or sap-net) and a unique *cc-id*. The *ccag-id* identifies the group of CCAs that will be used for forwarding packets associated with the SAP. The path identifies the bandwidth control grouping used to manage CCA egress bandwidth. The pairing type helps the system identify which buffering resources will be used to manage egress queuing of packets. Finally, the *cc-id* is used to explicitly cross connect the SAP to another SAP or network IP interface configured with the same *cc-id*.

Services Commands

sap

Syntax **sap ccag**-*ccag-id*.{**a** | **b**}[.*sap-net* | .*sap-sap*]:*cc-id* [**create**]
no sap ccag-*ccag-id*.{**a** | **b**}[.*sap-net* | .*sap-sap*]:*cc-id*

Context config>service>epipe

Description This command creates a cross connect SAP on the *ccag-id* referenced in the Epipe service. A CCAG SAP is identified by four items; the *ccag-id*, the CCAG path, the pairing type and the *cc-id*. An Epipe CCAG SAP supports all the available QoS, filtering and accounting features as an Epipe Dot1Q SAP.

The **no** form of the command removes a SAP from a service context. Once removed, all information and resources concerning the SAP is deleted from the system including the CCAG *cc-id* in use on the CCA path.

Parameters **ccag —** The ccag portion of the SAP identifier is required and specifies that the epipe SAP is of the CCAG type.

-*ccag-id* — The *ccag-id* portion of the SAP identifier is required and specifies which *ccag-id* on which the SAP must be created. The specified *ccag-id* must exist on the system or the SAP creation will fail. The leading dash must be included as a separator between **ccag** and the *ccag-id*.

Values -1 (dash 1) to -8 (dash 8)

Default **None**

.a | .b — The .a and .b portion of the CCAG SAP identifier is required and is used to define the CCA bandwidth path with will be associated with the SAP. The path association must be specified and .a and .b are mutually exclusive. The .a designation identifies the SAP as being on the Alpha path and the .b designation identifies the SAP as being on the Beta path. The paired SAP or network IP interface using the same *cc-id* as the SAP must be associated with the opposite path. The leading period must be included as a separator between the *ccag-id* and the path designator.

Values .a or .b

Default **None**

.*sap-net* — The .sap-net portion of the CCAG SAP identifier specifies that the SAP is of the SAP-2-Net pairing type and is required when the *cc-id* is paired with a network IP interface. The pairing type .sap-net is mutually exclusive with pairing type .sap-sap. If .sap-net is not specified, .sap-sap is assumed and does not need to be included in the SAP identification. When specified, the leading period must be used as a separator between the path designator and the pairing type.

Values .sap-net or .sap-sap

Default **.sap-sap**

.*sap-sap* — The .sap-sap portion of the CCAG SAP identifier is mutually exclusive to .sap-sap and is used to define the pairing type as SAP-2-SAP. The .sap-sap pairing type is only used when the cross connect object sharing the same *cc-id* on the opposite path is a CCAG SAP. If the other cross connect object is a network IP interface, the pairing type must be defined as .sap-net. If

.sap-net is not specified, .sap-sap is assumed and does not need to be included in the SAP identification. When specified, the leading period must be used as a separator between the path designator and the pairing type.

Values .sap-net or .sap-sap

Default **.sap-sap**

:cc-id — The :*cc-id* portion of the CCAG SAP identifier is required and specifies the unique *cc-id* in use by the CCAG SAP and the cross connect object on the other path.

Values 1 to 4094

Default **None**

create — Explicitly indicates that the specified CCAG SAP is being created by the **sap** command. Handling the inclusion or exclusion state of the create keyword is dependent on the system environment variable create.

When the system environment variable create is enabled, the system requires the explicit use of the **create** keyword when creating objects such as SAPs. If the keyword is not included and the specified CCAG SAP has not already been created, an error will occur and the CLI will not change context to the specified CCAG SAP instance. This is designed to prevent the inadvertent creation of a CCAG SAP in the event where the wrong CCAG SAP identifier is specified during an attempt to edit an existing CCAG SAP. If the **create** keyword is specified, the CCAG SAP will be created if it does not already exist or if it does exist, the CLI context will change to the specified CCAG SAP.

When the system environment variable create is disabled (using the **no create** command), the system will not require the **create** keyword when creating a CCAG SAP. In the event that the **sap** command is issued with a CCAG SAP identifier that previously had not been created, that CCAG SAP will be created.

Once a CCAG SAP has been created, the **create** keyword is ignored when a **sap** command is executed with that CCAG SAP identifier and the CLI context will change to the specified CCAG SAP.

sap

Syntax **sap ccag**-*ccag-id*.{**a** | **b**}[.*sap-net* | .*sap-sap*]:*cc-id* [**create**]
no sap ccag-*ccag-id*.{**a** | **b**}[.*sap-net* | .*sap-sap*]:*cc-id*

Context config>service>vpls

Description This command creates a cross connect SAP on the *ccag-id* referenced in the VPLS service. A CCAG SAP is identified by four items; the *ccag-id*, the CCAG path, the pairing type and the *cc-id*. A VPLS CCAG SAP supports all the available QoS, filtering and accounting features as a VPLS Dot1Q SAP.

The **no** form of the command removes a SAP from a service context. Once removed, all information and resources concerning the SAP is deleted from the system including the CCAG *cc-id* in use on the CCA path.

Parameters **ccag —** The ccag portion of the SAP identifier is required and specifies that the vpls SAP is of the CCAG type.

-ccag-id — Specifies which *ccag-id* on which the SAP must be created. The specified *ccag-id* must exist on the system or the SAP creation will fail. The leading dash must be included as a separator between **ccag** and the *ccag-id*.

Values -1 (dash 1) to -8 (dash 8)

Default **None**

.a | .b — The **.a** and **.b** portion of the CCAG SAP identifier is required and is used to define the CCA bandwidth path with will be associated with the SAP. The path association must be specified and .a and .b are mutually exclusive. The .a designation identifies the SAP as being on the Alpha path and the .b designation identifies the SAP as being on the Beta path. The paired SAP or network IP interface using the same *cc-id* as the SAP must be associated with the opposite path. The leading period must be included as a separator between the *ccag-id* and the path designator.

Values .a or .b

Default **None**

.*sap-net* — The .sap-net portion of the CCAG SAP identifier specifies that the SAP is of the SAP-2-Net pairing type and is required when the *cc-id* is paired with a network IP interface. The pairing type .sap-net is mutually exclusive with pairing type .sap-sap. If .sap-net is not specified, .sap-sap is assumed and does not need to be included in the SAP identification. When specified, the leading period must be used as a separator between the path designator and the pairing type.

Values .sap-net or .sap-sap

Default **.sap-sap**

.*sap-sap* — The .sap-sap portion of the CCAG SAP identifier is mutually exclusive to .sap-sap and is used to define the pairing type as SAP-2-SAP. The .sap-sap pairing type is only used when the cross connect object sharing the same *cc-id* on the opposite path is a CCAG SAP. If the other cross connect object is a network IP interface, the pairing type must be defined as .sap-net. If .sap-net is not specified, .sap-sap is assumed and does not need to be included in the SAP identification. When specified, the leading period must be used as a separator between the path designator and the pairing type.

Values .sap-net or .sap-sap

Default **.sap-sap**

:cc-id — The *:cc-id* portion of the CCAG SAP identifier is required and specifies the unique *cc-id* in use by the CCAG SAP and the cross connect object on the other path.

Values 1 to 4094

Default **None**

create — Explicitly indicates that the specified CCAG SAP is being created by the **sap** command. Handling the inclusion or exclusion state of the create keyword is dependent on the system environment variable create.

When the system environment variable create is enabled, the system requires the explicit use of the **create** keyword when creating objects such as SAPs. If the keyword is not included and the specified CCAG SAP has not already been created, an error will occur and the CLI will not change context to the specified CCAG SAP instance. This is designed to prevent the inadvertent creation of a CCAG SAP in the event where the wrong CCAG SAP identifier is specified during an attempt to edit an existing CCAG SAP. If the **create** keyword is specified, the CCAG SAP

will be created if it does not already exist or if it does exist, the CLI context will change to the specified CCAG SAP.

When the system environment variable create is disabled (using the **no create** command), the system will not require the **create** keyword when creating a CCAG SAP. In the event that the **sap** command is issued with a CCAG SAP identifier that previously had not been created, that CCAG SAP will be created.

Once a CCAG SAP has been created, the **create** keyword is ignored when a **sap** command is executed with that CCAG SAP identifier and the CLI context will change to the specified CCAG SAP.

sap

Syntax **sap ccag**-*ccag-id*.{**a** | **b**}[.*sap-net* | .*sap-sap*]:*cc-id* [**create**]
no sap ccag-*ccag-id*.{**a** | **b**}[.*sap-net* | .*sap-sap*]:*cc-id*

Context config>service>ies>interface

Description This command creates a cross connect SAP on the *ccag-id* referenced in the IES service. A CCAG SAP is identified by four items; the *ccag-id*, the CCAG path, the pairing type and the *cc-id*. A CCAG SAP on an IES IP interface supports all the available QoS, filtering and accounting features as an IES IP interface Dot1Q SAP.

The **no** form of the command removes a SAP from the IES service IP interface context. Once removed, all information and resources concerning the SAP is deleted from the system including the CCAG *cc-id* in use on the CCA path.

Parameters **ccag —** The **ccag** portion of the SAP identifier is required and specifies that the ies SAP is of the CCAG type.

-*ccag-id* — The *ccag-id* portion of the SAP identifier is required and specifies which *ccag-id* on which the SAP must be created. The specified *ccag-id* must exist on the system or the SAP creation will fail. The leading dash must be included as a separator between ccag and the *ccag-id*.

Values -1 (dash 1) to -8 (dash 8)

Default **None**

.a | .b — The .a and .b portion of the CCAG SAP identifier is required and is used to define the CCA bandwidth path with will be associated with the SAP. The path association must be specified and .a and .b are mutually exclusive. The .a designation identifies the SAP as being on the Alpha path and the .b designation identifies the SAP as being on the Beta path. The paired SAP or network IP interface using the same *cc-id* as the SAP must be associated with the opposite path. The leading period must be included as a separator between the *ccag-id* and the path designator.

Values .a or .b

Default **None**

.*sap-sap* — The .sap-sap portion of the CCAG SAP identifier is optional and is used to explicitly define the pairing type as SAP-2-SAP. The .sap-sap pairing type is only used when the cross connect object sharing the same *cc-id* on the opposite path is a CCAG SAP. A cross connection between an IES CCAG SAP and a network IP interface is not currently allowed. If .sap-sap is not

specified, .sap-sap is assumed and does not need to be included in the SAP identification. When specified, the leading period must be used as a separator between the path designator and the pairing type.

Default **.sap-sap**

:cc-id — The *:cc-id* portion of the CCAG SAP identifier is required and specifies the unique *cc-id* in use by the CCAG SAP and the cross connect object on the other path.

Values 1 to 4094

Default **None**

create — Explicitly indicates that the specified CCAG SAP is being created by the **sap** command. Handling the inclusion or exclusion state of the create keyword is dependent on the system environment variable create.

When the system environment variable create is enabled, the system requires the explicit use of the **create** keyword when creating objects such as SAPs. If the keyword is not included and the specified CCAG SAP has not already been created, an error will occur and the CLI will not change context to the specified CCAG SAP instance. This is designed to prevent the inadvertent creation of a CCAG SAP in the event where the wrong CCAG SAP identifier is specified during an attempt to edit an existing CCAG SAP. If the **create** keyword is specified, the CCAG SAP will be created if it does not already exist or if it does exist, the CLI context will change to the specified CCAG SAP.

When the system environment variable create is disabled (using the **no create** command), the system will not require the **create** keyword when creating a CCAG SAP. In the event that the **sap** command is issued with a CCAG SAP identifier that previously had not been created, that CCAG SAP will be created.

Once a CCAG SAP has been created, the **create** keyword is ignored when a **sap** command is executed with that CCAG SAP identifier and the CLI context will change to the specified CCAG SAP.

sap

Syntax **sap ccag**-*ccag-id*.{**a** | **b**}[.*sap-net* | .*sap-sap*]:*cc-id* [**create**]
no sap ccag-*ccag-id*.{**a** | **b**}[.*sap-net* | .*sap-sap*]:*cc-id*

Context config>service>vprn>interface

Description This command creates a cross connect SAP on the *ccag-id* referenced in the VPRN service. A CCAG SAP is identified by four items; the *ccag-id*, the CCAG path, the pairing type and the *cc-id*. A CCAG SAP on a VPRN IP interface supports all the available QoS, filtering and accounting features as a VPRN IP interface Dot1Q SAP.

The **no** form of the command removes a SAP from the VPRN service IP interface context. Once removed, all information and resources concerning the SAP is deleted from the system including the CCAG *cc-id* in use on the CCA path.

Parameters **ccag —** The ccag portion of the SAP identifier is required and specifies that the vprn SAP is of the CCAG type.

-ccag-id — Specifies which *ccag-id* on which the SAP must be created. The specified *ccag-id* must exist on the system or the SAP creation will fail. The leading dash must be included as a separator between **ccag** and the *ccag-id*.

Values -1 (dash 1) to -8 (dash 8)

Default **None**

.a | .b — The .**a** and **.b** portion of the CCAG SAP identifier is required and is used to define the CCA bandwidth path with will be associated with the SAP. The path association must be specified and .**a** and .**b** are mutually exclusive. The .a designation identifies the SAP as being on the alpha path and the **.b** designation identifies the SAP as being on the beta path. The paired SAP or network IP interface using the same *cc-id* as the SAP must be associated with the opposite path. The leading period must be included as a separator between the *ccag-id* and the path designator.

Values .a or .b

Default **None**

.*sap-net* — Specifies that the SAP is of the SAP-2-Net pairing type and is required when the *cc-id* is paired with a network IP interface. The pairing type .sap-net is mutually exclusive with pairing type .*sap-sap*. If .*sap-net* is not specified, .*sap-sap* is assumed and does not need to be included in the SAP identification. When specified, the leading period must be used as a separator between the path designator and the pairing type.

Values .sap-net or .sap-sap

Default **.sap-sap**

.*sap-sap* — The .*sap-sap* portion of the CCAG SAP identifier is mutually exclusive to .*sap-net* and is used to define the pairing type as SAP-2-SAP. The .*sap-sap* pairing type is only used when the cross connect object sharing the same *cc-id* on the opposite path is a CCAG SAP. If the other cross connect object is a network IP interface, the pairing type must be defined as .*sap-net*. If .*sap-net* is not specified, .*sap-sap* is assumed and does not need to be included in the SAP identification. When specified, the leading period must be used as a separator between the path designator and the pairing type.

Values .sap-net or .sap-sap

Default **.sap-sap**

:cc-id — The :*cc-id* portion of the CCAG SAP identifier is required and specifies the unique *cc-id* in use by the CCAG SAP and the cross connect object on the other path.

Values 1 to 4094

Default **None**

create — Explicitly indicates that the specified CCAG SAP is being created by the **sap** command. Handling the inclusion or exclusion state of the **create** keyword is dependent on the system environment variable create.

When the system environment variable create is enabled, the system requires the explicit use of the **create** keyword when creating objects such as SAPs. If the keyword is not included and the specified CCAG SAP has not already been created, an error will occur and the CLI will not change context to the specified CCAG SAP instance. This is designed to prevent the inadvertent creation of a CCAG SAP in the event where the wrong CCAG SAP identifier is specified during an attempt to edit an existing CCAG SAP. If the create keyword is specified, the CCAG SAP will

be created if it does not already exist or if it does exist, the CLI context will change to the specified CCAG SAP.

When the system environment variable create is disabled (using the no create command), the system will not require the create keyword when creating a CCAG SAP. In the event that the sap command is issued with a CCAG SAP identifier that previously had not been created, that CCAG SAP will be created.

Once a CCAG SAP has been created, the **create** keyword is ignored when a **sap** command is executed with that CCAG SAP identifier and the CLI context will change to the specified CCAG SAP.

Mirror Services

In This Chapter

This chapter provides information to configure mirroring.

Topics in this chapter include:

Service Mirroring

When troubleshooting complex operational problems, customer packets can be examined as they traverse the network. One way to accomplish this is with an overlay of network analyzers established at multiple PoPs, together with skilled technicians to operate them to decode the data provided. This method of traffic mirroring often requires setting up complex filters in multiple switches and/or routers. These, at best, are only able to mirror from one port to another on the same device.

Alcatel-Lucent's Service Mirroring extends and integrates these capabilities into the network and provides significant operational benefits. Each 7750 SR can mirror packets from a specific service to any destination point in the network, regardless of interface type or speed.

This capability also extends beyond troubleshooting services. Telephone companies have the ability to obtain itemized calling records and wire-taps where legally required by investigating authorities. The process can be very complex and costly to carry out on data networks. Service Mirroring greatly simplifies these tasks, as well as reduces costs through centralization of analysis tools and skilled technicians.

Alcatel-Lucent's 7750 SR routers support service-based mirroring. While some Layer 3 switches and routers can mirror on a per-port basis within the device, Alcatel-Lucent 7750 SR routers can mirror on an n-to-1 unidirectional service basis and re-encapsulate the mirrored data for transport through the core network to another location, using either IP or MPLS tunneling as required (Figure 61).

Original packets are forwarded while a copy is sent out the mirrored port to the mirroring (destination) port. Service mirroring allows an operator to see the actual traffic on a customer's service with a sniffer sitting in a central location. In many cases, this reduces the need for a separate, costly overlay sniffer network.

The mirrored frame size that is to be transmitted to the mirror destination can be explicitly configured by using slicing features. This enables mirroring only the parts needed for analysis. For example, only the headers can be copied for analysis, protecting the integrity and security of customer data, or conversely, copying the full packet, including customer data.

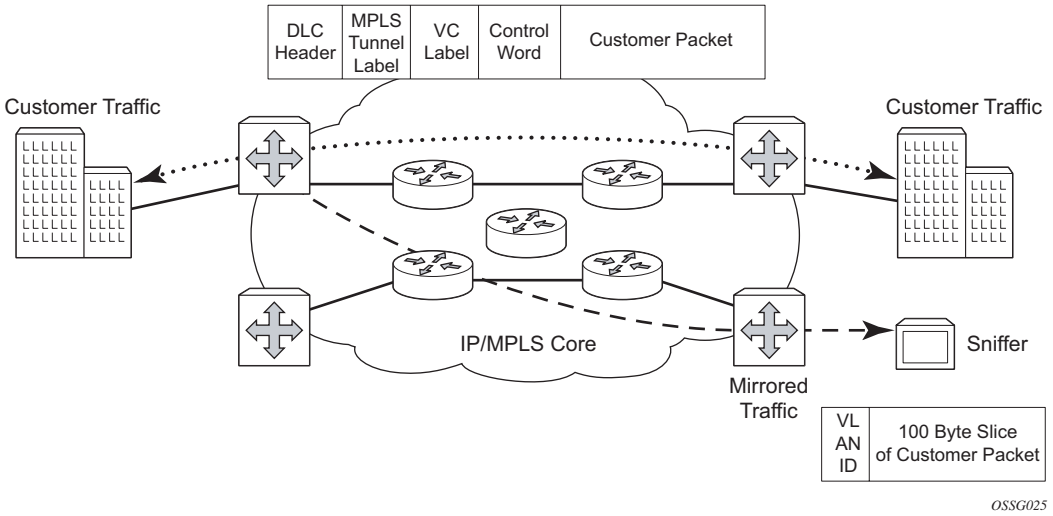

Figure 1: Service Mirroring

Mirror Implementation

Mirroring can be implemented on ingress or egress service access points (SAPs) or ingress and egress network interfaces. The Flexible Fast Path processing complexes preserve the ingress packet throughout the forwarding and mirroring process, making incremental packet changes on a separate copy.

Alcatel-Lucent's implementation of packet mirroring is based on two assumptions:

- Ingress and egress packets are mirrored as they appear on the wire. This is important for troubleshooting encapsulation and protocol issues.
  - → When mirroring at ingress, the Flexible Fast Path network processor array (NPA) sends an exact copy of the original ingress packet to the mirror destination while normal forwarding proceeds on the original packet.
  - → When mirroring is at egress, the NPA performs normal packet handling on the egress packet, encapsulating it for the destination interface. A copy of the forwarded packet (as seen on the wire) is forwarded to the mirror destination.
- Mirroring must support tunnel destinations.
  - → Remote destinations are reached by encapsulating the ingress or egress packet within an SDP, like the traffic for distributed VPN connectivity services. At the remote destination, the tunnel encapsulation is removed and the packet is forwarded out a local SAP.

Mirror Source and Destinations

Mirror sources and destinations have the following characteristics:

- They can be on the same 7750 SR router (local) or on two different routers (remote).
- Mirror destinations can terminate on egress virtual ports which allows multiple mirror destinations to send to the same packet decode device, delimited by IEEE 802.1Q (referred to as dot1q) tags. This is helpful when troubleshooting a multi-port issue within the network.

  When multiple mirror destinations terminate on the same egress port, the individual dot1q tags can provide a DTE/DCE separation between the mirror sources.
- Packets ingressing a port can have a mirror destination separate from packets egressing another or the same port (the ports can be on separate nodes).
- A total of 255 mirror destinations are supported (local and/or remote), per chassis.

Local and Remote Mirroring

Mirrored frames can be copied and sent to a specific local destination or service on the 7750 SR router (local mirroring) or copies can be encapsulated and sent to a different 7750 SR router (remote mirroring). This functionality allows network operators to centralize not only network analyzer (sniffer) resources, but also the technical staff who operate them.

The 7750 SR allows multiple concurrent mirroring sessions so traffic from more than one ingress mirror source can be mirrored to the same or different egress mirror destinations.

Remote mirroring uses a service distribution path (SDP) which acts as a logical way of directing traffic from one 7750 SR-Series router to another through a uni-directional (one-way) service tunnel. The SDP terminates at the far-end 7750 SR which directs packets to the correct destination on that device.

The SDP configuration from the mirrored device to a far-end 7750 SR requires a return path SDP from the far-end 7750 SR back to the mirrored router. Each device must have an SDP defined for every remote router to which it wants to provide mirroring services. SDPs must be created first, before services can be configured.

Slicing

A further service mirroring refinement is 'slicing' which copies a specified packet size of each frame. This is useful to monitor network usage without having to copy the actual data. Slicing enables mirroring larger frames than the destination packet decode equipment can handle. It also allows conservation of mirroring resources by limiting the size of the stream of packet through the 7750 SR and the core network.

When a mirror `slice-size` is defined, a threshold that truncates a mirrored frame to a specific size is created. For example, if the value of 256 bytes is defined, up to the first 256 bytes of the frame are transmitted to the mirror destination. The original frame is not affected by the truncation. Mirrored frames, most likely, will grow larger as encapsulations are added when packets are transmitted through the network core or out the mirror destination SAP to the packet/protocol decode equipment.

The transmission of a sliced or non-sliced frame is also dependent on the mirror destination SDP path MTU and/or the mirror destination SAP physical MTU. Packets that require a larger MTU than the mirroring destination supports are discarded if the defined slice size does not truncate the packet to an acceptable size.

Mirroring Performance

Replication of mirrored packets can, typically, affect performance and should be used carefully. Alcatel-Lucent 7750 SR routers minimize the impact of mirroring on performance by taking advantage of its distributed Flexible Fast Path technology. Flexible Fast Path forwarding allows efficient mirror service scaling and, at the same time, allows a large amount of data to be mirrored with minimal performance impact. When a mirror destination is configured, the packet slice option can truncate mirrored packets to the destination, which minimizes replication and tunneling overhead. The mirroring architecture also supports mirror rate limiting both at the ingress and egress Flexible Fast Path NPA. This rate limiting is accomplished though a shaping queue and is set according to the maximum amount of mirroring desired.

Mirroring can be performed based on the following criteria:

- Port
- SAP
- MAC filter
- IP filter
- Ingress label

ATM Mirroring

ATM mirror functionality allows 7750 SR users to mirror AAL5 packets from a source ATM SAP to a destination ATM SAP connected locally or remotely. This functionality can be used to monitor the ATM traffic on a particular ATM SAP. In both the local and remote scenarios the source and destination SAPs must be of ATM SAP type.

All ingress and egress AAL5 traffic at the source ATM SAP is duplicated and sent toward the destination ATM SAP. Mirroring the ingress traffic only, egress traffic only, or both, can be configured. ATM OAM traffic is not mirrored toward the destination ATM SAP.

ATM mirroring is applicable to the following services using an ATM SAP:

- Layer 3: IES and VPRN
- Layer 2: Apipe (sdu-type only), Ipipe, EPipe, VPLS

IP filters used as a mirror source are supported on ATM SAPs based on the IP filter applicability for different services.

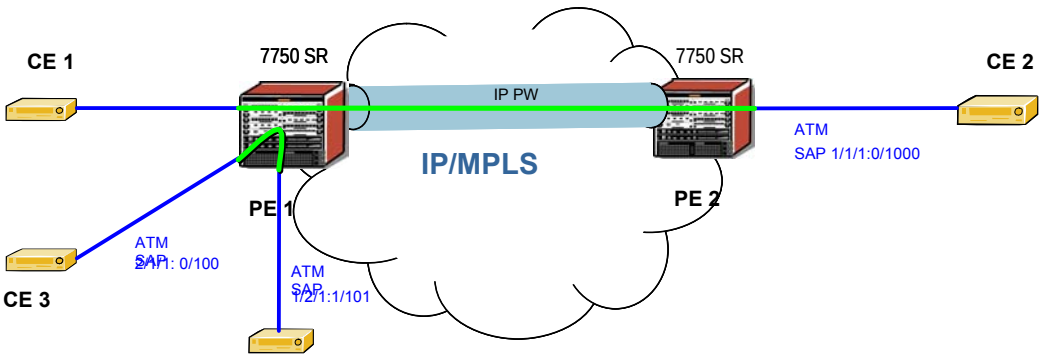

Figure 2: Example of an ATM Mirror Service

In Figure 62, CE 3 is connected to PE1 on ATM SAP 2/1/1/:0/100 as part of an IES service. The traffic on ATM SAP 2/1/1/:0/100 is mirrored locally to CE4 device through ATM SAP 1/2/1:1/101. In this scenario, all AAL5 packets arriving at SAP 2/1/1/:0/100 are duplicated and send towards ATM SAP 1/2/1:1/101.

In the case where the destination ATM SAP is on a remote node PE2, then the AAL5 traffic arriving at ATM SAP 2/1/1/:0/100 is duplicated and sent across the IP/MPLS network to PE2. At PE2 the traffic is forwarded to ATM SAP 1/1/1:0/1000 towards the ATM traffic monitoring device.

Mirroring Configuration

Configuring mirroring is similar to creating a uni-direction service. Mirroring requires the configuration of:

- Mirror source - the traffic on a specific point(s) to mirror.
- Mirror destination - the location to send the mirrored traffic, where the sniffer will be located.

Figure 63 depicts a local mirror service configured on ALA-A.

- Port 2/1/2 is specified as the source. Mirrored traffic ingressing and egressing this port will be sent to port 2/1/3.
- SAP 2/1/3 is specified as the destination. The sniffer is physically connected to this port. Mirrored traffic ingressing and egressing port 2/1/2 is sent here. SAP, encapsulation requirements, packet slicing, and mirror classification parameters are configured. SDPs are not used in local mirroring.

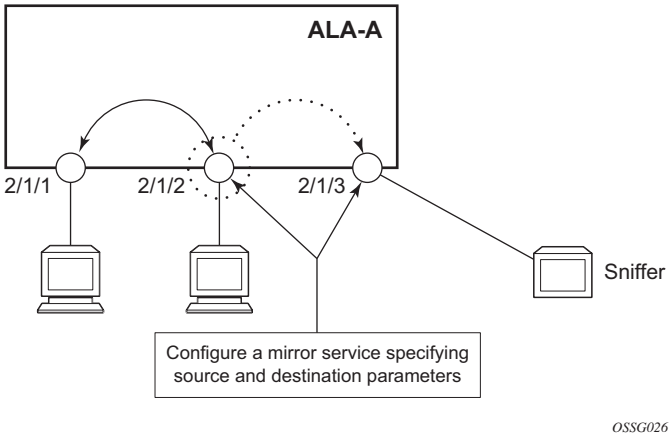

Figure 3: Local Mirroring Example

Figure 64 depicts a remote mirror service configured as ALA B as the mirror source and ALA A as the mirror destination. Mirrored traffic ingressing and egressing port 5/2/1 (the source) on ALA B is handled the following ways:

- Port 5/2/1 is specified as the *mirror source* port. Parameters are defined to select specific traffic ingressing and egressing this port.

  Destination parameters are defined to specify where the mirrored traffic will be sent. In this case, mirrored traffic will be sent to a SAP configured as part of the mirror service on port 3/1/3 on ALA A (the *mirror destination*).

  ALA A decodes the service ID and sends the traffic out of port 3/1/3.
  The sniffer is physically connected to this port (3/1/3). SAP, encapsulation requirements, packet slicing, and mirror classification parameters are configured in the destination parameters.

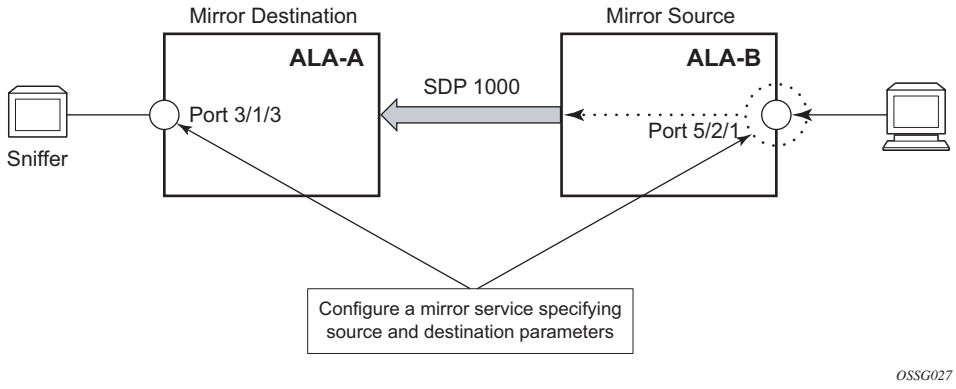

Figure 4: Remote Mirroring Example

Mirror Configuration Process Overview

Figure 65 displays the process to provision basic mirroring parameters.

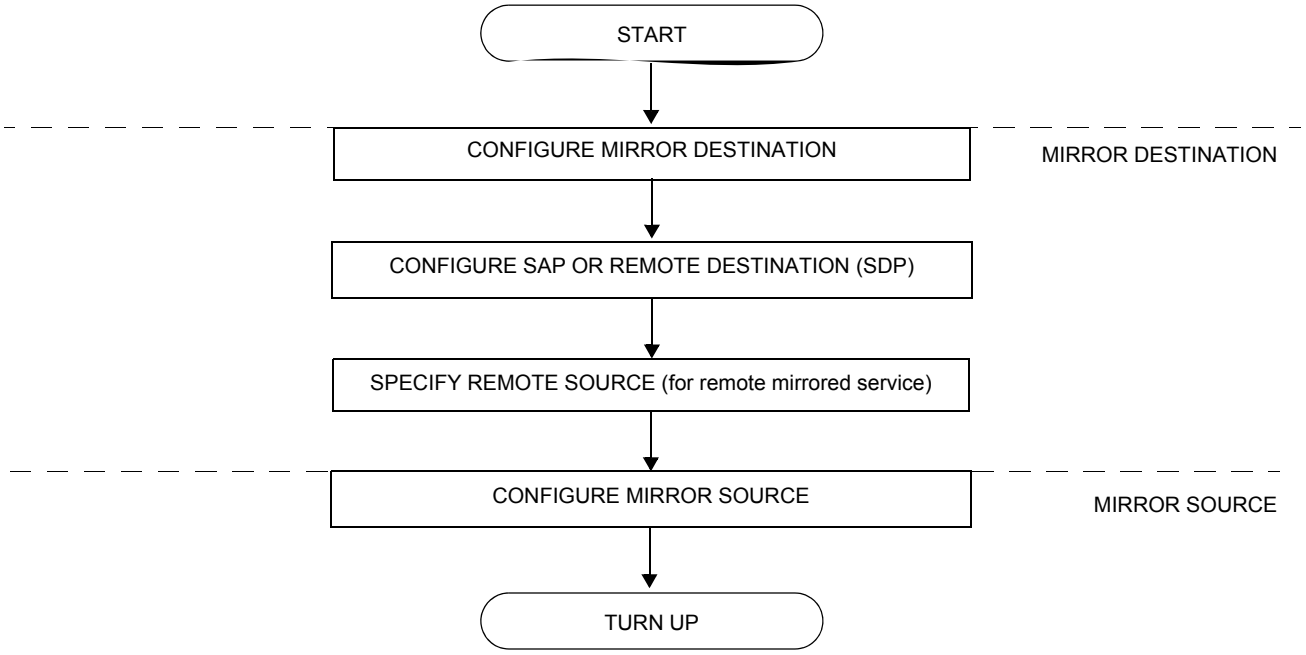

Figure 5: Mirror Configuration and Implementation Flow

Service Mirror Configuration Components

Figure 66 displays the major components to configure service mirroring.

```
CONFIGURE
    MIRROR DESTINATION SERVICE-ID
        SAP
        SDP (for remote mirrored service)
        REMOTE SOURCE (for remote mirrored service)

DEBUG
    MIRROR SOURCE SERVICE-ID
        PORT
        SAP
        IP FILTER
        MAC FILTER
        INGRESS LABEL
```

Figure 6: Service Mirroring Configuration Components

- Mirror destination — Sets up a service which allows the mirrored packets to be directed locally or over the core of the network and have a far end 7750 SR decode the mirror encapsulation. The service ID must match in the mirror-destination and the mirror-source context.
- SAP (mirror destination) — Creates a service access point (SAP), which defines the port and encapsulation parameters to which the mirrored source packets are sent. The sniffer is physically connected to this port.
- SDP — For remote mirrored service. Binds an existing (mirror) service distribution path (SDP) to the mirror destination service ID to transport the source mirrored traffic to the destination.
- Remote source — For remote mirrored services. Specifies the remote (source) device allowed to mirror traffic to this device for mirror service egress.
- Mirror source — Configures packet mirroring match criteria for a mirror destination service. The same mirror destination service ID and the mirror source service ID must be configured.
- Port — A packet mirroring option which defines ingress and/or egress traffic monitoring by port.
- SAP (mirror source) — A packet mirroring option which defines ingress and/or egress traffic monitoring by SAP defined by the SAP ID.

- IP filter — A packet mirroring option which specifies that packets matching the IP filter are mirrored to a mirror destination.
- MAC filter — A packet mirroring option which specifies that packets matching the MAC filter are mirrored to a mirror destination.
- Ingress label — A packet mirroring option which defines packets with a specific MPLS label to a mirror destination.

Configuration Notes

This section describes mirroring configuration caveats.

General

- Up to 255 mirroring service IDs may be created within a single system.
- A mirrored source can only have one destination.
- The destination mirroring service IDs and service parameters are persistent between router (re)boots and are included in the configuration saves.

  The source packet mirroring enabling criteria defined in debug mirror mirror-source commands are not preserved in configuration saves.
- Physical layer problems such as collisions, jabbers, etc., are not mirrored. Typically, only complete packets are mirrored. Complete stats are available on the interface for these physical layer problems.
- Starting and shutting down mirroring:

  Mirror destinations:

  → The default state for a mirror destination service ID is `shutdown`. You must issue a `no shutdown` command to enable the feature.

  → When a mirror destination service ID is shutdown, mirrored packets associated with the service ID are not accepted from its mirror source or remote source 7750 SR router. The associated mirror source is put into an operationally down mode. Mirrored packets are not transmitted out the SAP or SDP. Each mirrored packet is silently discarded. If the mirror destination is a SAP, the SAP's discard counters are incremented.

  → Issuing the `shutdown` command causes the mirror destination service or its mirror source to be put into an administratively down state. Mirror destination service IDs must be shut down first in order to delete a service ID, SAP, or SDP association from the system.

  Mirror sources:

  → The default state for a mirror source for a given mirror-dest service ID is `no shutdown`. You must enter a `shutdown` command to deactivate (disable) mirroring from that mirror-source.

  → Mirror sources do not need to be shutdown to remove them from the system. When a mirror source is shutdown, mirroring is terminated for all sources defined locally for the mirror destination service ID.

Reference Sources

For information on supported IETF drafts and standards, as well as standard and proprietary MIBS, refer to Standards and Protocol Support on page 1471.

Configuring Service Mirroring with CLI

This section provides information about service mirroring

Topics in this section include:

Mirror Configuration Overview

7750 SR mirroring can be organized in the following logical entities:

- The *mirror source* is defined as the location where ingress or egress traffic specific to a port, SAP, MAC or IP filter, or ingress label or a particular service is to be mirrored (copied). The original frames are not altered or affected in any way.
- An SDP is used to define the *mirror destination* on the source router to point to a remote destination (another router).
- A SAP is defined in local and remote mirror services as the *mirror destination* to where the mirrored packets are sent.

Defining Mirrored Traffic

In some scenarios, like using VPN services or when multiple services are configured on the same port, specifying the port does not provide sufficient resolution to separate traffic. In Alcatel-Lucent's implementation of mirroring, multiple source mirroring parameters can be specified to further identify traffic.

Mirroring of packets matching specific filter entries in an IP or MAC filter can be applied to refine what traffic is mirrored to flows of traffic within a service. The IP criteria can be combinations of:

- Source IP address/mask
- Destination IP address/mask
- IP Protocol value
- Source port value/range (for example, UDP or TCP port)
- Destination port value/range (for example, UDP or TCP port)
- DiffServ Code Point (DSCP) value
- IP fragments
- IP option value/mask
- Single or multiple IP option fields present
- IP option fields present
- TCP ACK set/reset
- TCP SYN set/reset
- ICMP code
- ICMP type
- SAP ingress/egress labels

The MAC criteria can be combinations of:

- IEEE 802.1p value/mask
- Source MAC address/mask
- Destination MAC address/mask
- Ethernet Type II Ethernet type value
- Ethernet 802.2 LLC DSAP value/mask
- Ethernet 802.2 LLC SSAP value/mask
- IEEE 802.3 LLC SNAP Ethernet Frame OUI zero/non-zero value
- IEEE 802.3 LLC SNAP Ethernet Frame PID value
- SAP ingress/egress labels

Mirror CLI Command Structure

Figure 67 displays the CLI command structure to configure mirroring parameters. The mirror destination commands are configured under the `config>mirror` context. The mirror source commands are configured under the `debug>mirror-source` context. The mirror destination configuration includes configuring the 7750 SR router with the sniffer to accept remote mirror sources and configuring the mirror service on the mirror source 7750 SR router for remote mirroring.

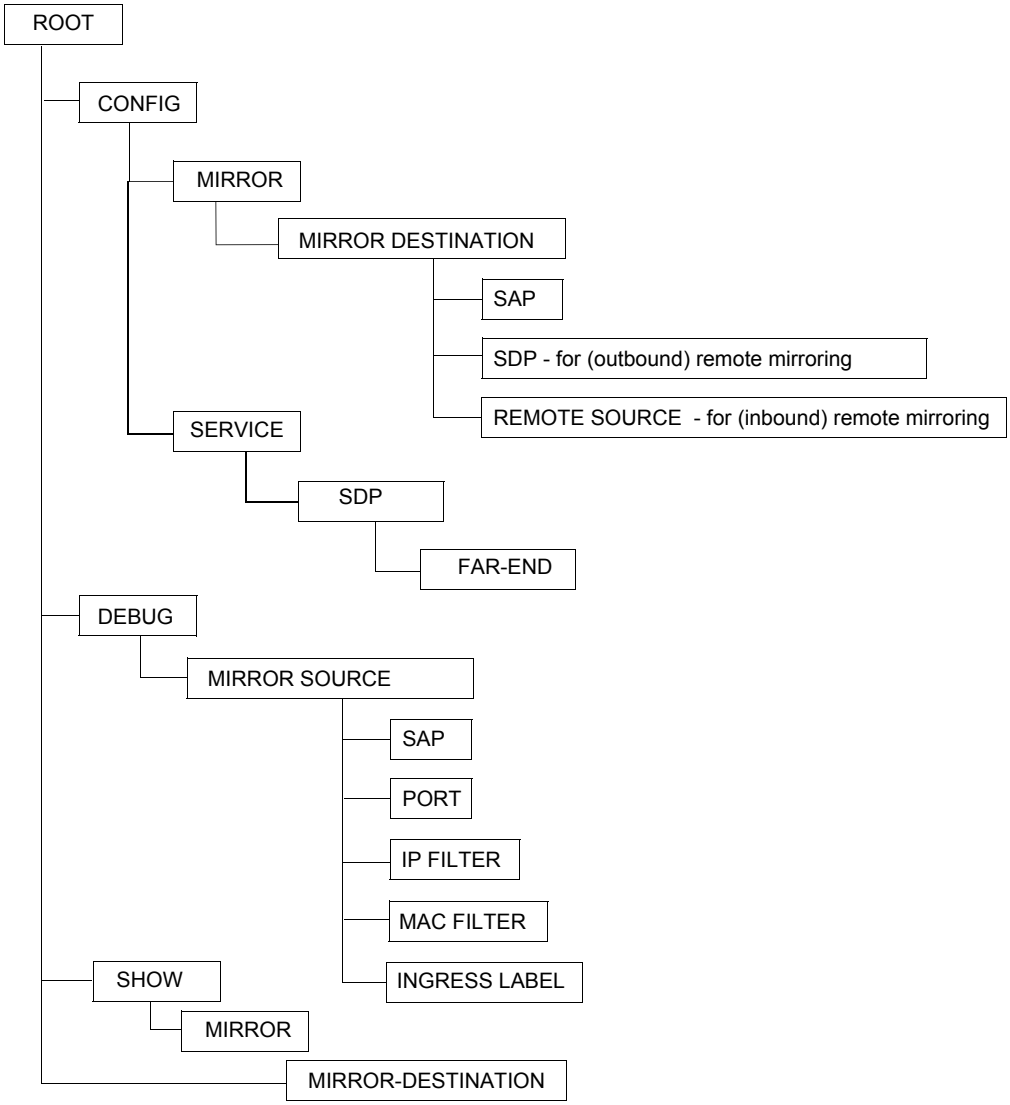

Figure 1: Mirror CLI Configuration Context

List of Commands

Table 41 lists all the configuration commands to configure mirroring parameters, indicating the configuration level at which each command is implemented with a short command description. The command list is organized in the following task-oriented manner:

- Configure mirror destination parameters
- Configure mirror source parameters
- Configure an SDP

Table 1: CLI Commands to Configure Mirroring Parameters

Command	Description	Page
Configure mirror destination parameters		
config>mirror		
mirror-dest	Creates a context for configuring a packet mirroring service. The mirroring instance is set up like a service to allow the mirrored packets to be directed over the core of the network and have a far end 7750 SR decode the mirror encapsulation.	1346
service-id	A unique service identification value which identifies this service in the service domain.	1347
description	A text description providing details of the mirroring instance.	1343
sap	Creates a Service Access Point (SAP) within a mirror destination service. A sniffer can be attached to this port.	1348
sdp	Binds an existing Service Distribution Point (SDP) to the mirror destination service ID. SDPs are used only for remote mirroring.	1351
remote-source	The context to configure remote devices allowed to mirror traffic to this device for mirror service egress.	1347
fc	Specifies a forwarding class for all mirrored packets transmitted to the destination SAP or SDP overriding the default forwarding class. All packets are sent with the same class of service to minimize out of sequence issues.	1346
far-end	Defines a remote source 7750 SR that may send mirrored packets to this 7750 SR for handling by the mirror destination service ID.	1345
egress	Configure SAP egress policies.	1351
qos	Configure an egress QoS policy for the SAP.	1351

Table 1: CLI Commands to Configure Mirroring Parameters (Continued)

Command	Description	Page
slice-size	Enables mirrored frame truncation and configures the maximum portion of the mirrored frame that will be transmitted to the mirror destination.	1353
no shutdown	Administratively enables the mirroring instance.	1343
Configure mirror source parameters		
debug>mirror		
mirror-source	Configures packet mirroring match criteria for a mirror service. This configuration is not saved.	1357
service-id	This is the mirror service ID which identifies the service. This service must be created first in the *mirror destination* context.	1357
port	Enables mirroring of traffic ingressing or egressing a port (Ethernet port, SONET/SDH or TDM channel, Link Aggregation Group (LAG)).	1358
sap	Enables mirroring of traffic ingressing or egressing a SAP.	1359
ip-filter	Enables mirroring of packets matching specific filter entries in an IP filter.	1354
ingress-label	Enables mirroring of ingress MPLS frames based on the top-of-stack MPLS label. Multiple labels can be defined simultaneously. The ingress label must be active in the source router.	1354
mac-filter	Enables mirroring of packets matching specific filter entries in a MAC filter.	1356
no shutdown	Administratively enables the instance for source mirroring.	1343
Configure an SDP		
config>service		
sdp	Creates a Service Distribution Point (SDP).	90
far-end	Configures the system IP address of the far-end destination 7750 SR for the SDP that is terminating services.	93
lsp	Creates associations between one or more label switched paths (LSPs) and an MPLS SDP. LSPs are configured in the config>router>mpls context.	94
path-mtu	Configures the Maximum Transmission Unit (MTU) in bytes that the SDP can transmit to the far-end 7750 SR without packet dropping the SDP-type default path-mtu.	96
no shutdown	Administratively enables the SDP.	79
keep-alive	Configures SDP connectivity monitoring keepalive messages for the SDP ID.	97
hello-time	Configures the time period between SDP keepalive messages on the SDP-ID for the SDP connectivity monitoring messages.	98

Table 1: CLI Commands to Configure Mirroring Parameters (Continued)

Command	Description	Page
`hold-down-time`	Configures the minimum time period the SDP will remain in the operationally down state in response to SDP keepalive monitoring.	98
`max-drop-count`	Configures the number of consecutive SDP keepalive failed request attempts or remote replies that can be missed after which the SDP is operationally downed.	98
`message-length`	Configures the size of SDP monitoring keepalive request messages.	99
`timeout`	Configures the time interval that the SDP waits before tearing down the session.	99
`no shutdown`	Administratively enables the keepalive messages.	79

Basic Mirroring Configuration

Destination mirroring parameters must include at least:

- A mirror destination ID (same as the mirror source service ID).
- A mirror destination SAP *or* SDP.

Mirror source parameters must include at least:

- A mirror service ID (same as the mirror destination service ID).
- At least one source type (port, SAP, ingress label, IP filter or MAC filter) specified.

The following example displays a sample configuration of a local mirrored service where the source and destinations are on the same device (ALA-A).

```
*A:ALA-A>config>mirror# info
----------------------------------------------
        mirror-dest 103 create
            sap 2/1/25:0 create
                egress
                    qos 1
                exit
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-A>config>mirror#
```

The following displays the *mirror source* configuration:

```
*A:ALA-A>debug>mirror-source# show debug mirror
debug
    mirror-source 103
        port 2/1/24 egress ingress
        no shutdown
    exit
exit
*A:ALA-A>debug>mirror-source# exit
```

The following example displays a sample configuration of a remote mirrored service where the source is a port on ALA-A and the destination a SAP is on ALA-B.

```
*A:ALA-A>config>mirror# info
----------------------------------------------
        mirror-dest 1000 create
            sdp 2 egr-svc-label 7000
            no shutdown
        exit
----------------------------------------------
*A:ALA-A>config>mirror# exit all
*A:ALA-A# show debug
debug
    mirror-source 1000
        port 2/1/2 egress ingress
        no shutdown
    exit
exit
*A:ALA-A#
```

```
*A:ALA-B>config>mirror# info
----------------------------------------------
        mirror-dest 1000 create
            remote-source
                far-end 10.10.10.104 ing-svc-label 7000
            exit
            sap 3/1/2:0 create
                egress
                    qos 1
                exit
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-B>config>mirror#
```

Mirror Classification Rules

Alcatel-Lucent's implementation of mirroring can be performed by configuring parameters to select network traffic according to any combination of the following entities:

- Port
- SAP
- MAC filter
- IP filter
- Ingress label

Port

The `port` command associates a port to a mirror source. The port is identified by the `port-id`. The following displays the port-id syntax:

port-id: slot/mda/port[.channel]
aps-*id* **aps**-*group-id*[.*channel*]
aps keyword
group-id 1 — 64

bundle-*type-slot/mda.bundle-num*
bundle keyword
type ima, ppp
bundle-num 1 — 128

ccag-*id* - **ccag**-*id.path-id*[*cc-type*]:*cc-id*
ccag keyword
id 1 — 8
path-id a, b
cc-type .sap-net, .net-sap
cc-id 0 — 4094
lag-id 1 — 64
egress keyword
ingress keyword

The defined port can be an ATM, Ethernet or Frame Relay port, a SONET/SDH path, a multilink bundle, a TDM channel, a Cross Connect Aggregation Group (CCAG), or a Link Aggregation Group (LAG) ID. If the port is a SONET/SDH or TDM channel, the channel ID must be specified to identify which channel is being mirrored. When a LAG ID is given as the port ID, mirroring is enabled on all ports making up the LAG.

Mirror sources can be ports in either access or network mode. Port mirroring is supported in the following combinations:

Table 2: Mirror Source Port Requirements

Port Type	Port Mode	Port Encap Type
faste/gige/xgige	access	dot1q, null, qinq
faste/gige/xgige	network	dot1q, null
SONET (clear channel)	network	ppp-auto
SONET (clear/deep channel)	access	bcp-null, bcp-dot1q, ipcp
TDM (clear/deep channel)	access	bcp-null, bcp-dot1q, ipcp

CLI Syntax: `debug>mirror-source# port {port-id|lag lag-id} {[egress][ingress]}`

Example: `*A:ALA-A>debug>mirror-source# port 2/2/2 ingress egress`

SAP

More than one SAP can be associated within a single `mirror-source`. Each SAP has its own ingress and egress parameter keywords to define which packets are mirrored to the mirror-dest service ID.

You cannot create a mirror destination SAP on a SONET/SDH or TDM channel. You can configure SONET/SDH or TDM channel SAPs as mirror sources.

The port requirements for mirror destination SAPs are as follows:

- The port type must be Fast Ethernet, Gigabit Ethernet, or XGigE.
- The port mode must be access.
- The port encapsulation types must be dot1q, null, or qinq.

CLI Syntax: `debug>mirror-source# sap` *`sap-id`* `{[egress] [ingress]}`

Example: `*A:ALA-A>debug>mirror-source# sap 2/1/4:100 ingress egress`

or `debug>mirror-source# port 2/2/1.sts12 ingress`

MAC filter

MAC filters are configured in the `config>filter>mac-filter` context. The `mac-filter` command causes all the packets matching the explicitly defined list of entry IDs to be mirrored to the mirror destination specified by the `service-id` of the mirror source.

CLI Syntax: `debug>mirror-source# mac-filter` *`mac-filter-id`* `entry` *`entry-id`* `[`*`entry-id`* `…]`

Example: `*A:ALA-2>debug>mirror-source# mac-filter 12 entry 15 20 25`

IP filter

IP filters are configured in the `config>filter>ip-filter` context. The `ip-filter` command causes all the packets matching the explicitly defined list of entry IDs to be mirrored to the mirror destination specified by the `service-id` of the mirror source.

Ingress mirrored packets are mirrored to the mirror destination prior to any ingress packet modifications. Egress mirrored packets are mirrored to the mirror destination after all egress packet modifications.

CLI Syntax: `debug>mirror-source# ip-filter` *`ip-filter-id`* `entry` *`entry-id`* `[`*`entry-id`* `…]`

Example: `*A:ALA-A>debug>mirror-source# ip-filter 1 entry 20`

NOTE: An IP filter cannot be applied to a mirror destination SAP.

Ingress label

The `ingress-label` command is used to mirror ingressing MPLS frames with the specified MPLS labels. The ingress label must be at the top of the label stack and can only be mirrored to a single mirror destination. If the same label is defined with multiple mirror destinations, an error is generated and the original mirror destination does not change. `Ingress-label` allows packets matching the ingress label to be duplicated (mirrored) and forwarded to the mirror destination. The ingress label has to be active before it can be used as mirror source criteria. If the ingress label is not used in the router, the mirror source will remove the ingress label automatically.

CLI Syntax: `debug>mirror-source# ingress-label` *`label`* `[`*`label`*`...]`

Example: `*A:ALA-A>debug>mirror-source# ingress-label 103000 1048575`

Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure both local and remote mirror services and provides the CLI command syntax. Note that local and remote *mirror source* and *mirror destination* components must be configured under the same service ID context.

Each local mirrored service (Figure 68) (within the same router) requires the following configurations:

1. Specify *mirror destination* (SAP, SDP).
2. Specify *mirror source* (port, SAP, SDP, IP filter, MAC filter, ingress label).

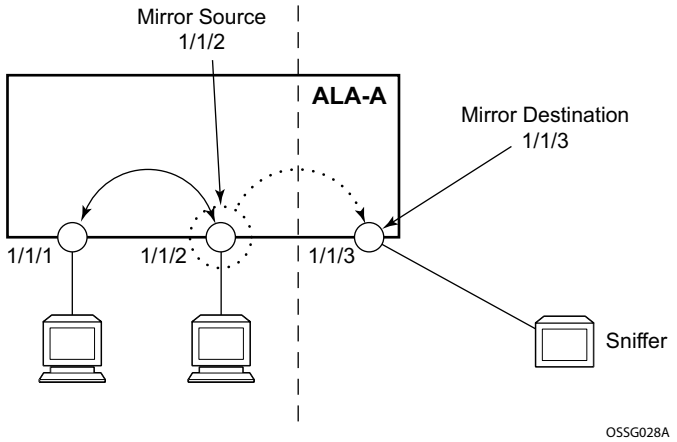

Figure 2: Local Mirrored Service Tasks

Each remote mirrored service (Figure 69) (across the network core) requires the following configurations:

1. Define the remote destination (SDP)
2. Identify the remote source (the device allowed to mirror traffic to this device)
3. Specify the mirror destination (SAP)
4. Specify mirror source (port, SAP, SDP, IP filter, MAC filter, ingress label)

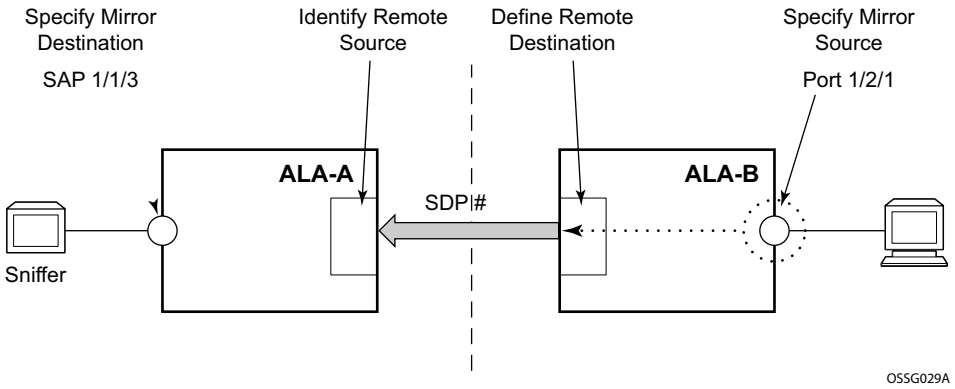

Figure 3: Remote Mirrored Service Tasks

Configuring a Local Mirror Service

To configure a local mirror service, the source and destinations must be located on the same router. Note that local mirror source and mirror destination components must be configured under the same service ID context.

The `mirror-source` commands are used as traffic selection criteria to identify traffic to be mirrored at the source. For example, use the `port port {port-id|lag lag-id} {[egress][ingress]}` and `ip-filter ip-filter-id entry entry-id [entry-id …]` commands to capture (mirror) traffic that matches specific IP filter entry criteria which is ingressing and egressing a specific port. A filter must be applied to the SAP/interface if only certain packets are to be mirrored.

Use the CLI syntax to configure one or more mirror source parameters:

The `mirror-dest` commands are used to specify where the mirrored traffic is to be sent, the forwarding class, and the size of the packet. Use the following CLI syntax to configure mirror destination parameters:

CLI Syntax:
```
config>mirror
 mirror-dest service-id [type {ether|frame-relay|ppp|atm-
 sdu}]
    description string
    fc fc-name
    sap sap-id
    slice-size bytes
    no shutdown
```

CLI Syntax:
```
debug# mirror-source service-id
 ip-filter ip-filter-id entry entry-id [entry-id …]
 ingress-label label [label …]
 mac-filter mac-filter-id entry entry-id [entry-id …]
 port {port-id|lag lag-id} {[egress][ingress]}
 sap sap-id {[egress][ingress]}
 no shutdown
```

The following output displays an example of a local mirrored service. On ALA-A, mirror service 103 is mirroring egress and ingress traffic matching IP filter 2, entry 1, on port 2/1/24 and sending the mirrored packets to SAP 2/1/25.

Example:
```
*A:ALA-A>config>mirror# mirror-dest 103 create
      config>mirror>mirror-dest$ sap 2/1/25:0 create
      config>mirror>mirror-dest>sap# exit
      config>mirror>mirror-dest# no shutdown
      config>mirror>mirror-dest# exit all
      debug>mirror-source 103
```

```
debug>mirror-source# port 2/1/24 egress ingress
debug>mirror-source# ip-filter 2 entry 1
```

The following displays the local mirroring configuration:

```
*A:ALA-A>config>mirror# info
----------------------------------------------
        mirror-dest 103 create
            sap 2/1/25:0 create
                egress
                    qos 1
                exit
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-A>config>mirror#
```

The following displays the debug mirroring information:

```
*A:ALA-A>debug>mirror-source# show debug mirror
debug
    mirror-source 103
        no shutdown
        port 2/1/24 egress ingress
        ip-filter 2 entry 1
    exit
exit
*A:ALA-A>debug>mirror-source# exit
```

Configuring SDPs

This section provides a brief overview of the tasks that must be performed to configure SDPs and provides the CLI commands. For more information about service configuration, refer to the *Subscriber Services* chapter.

Consider the following SDP characteristics:

- Configure either GRE or MPLS SDPs.
- Each distributed service must have an SDP defined for every remote SR to provide Epipe, VPLS, or mirrored services.
- A distributed service must be bound to an SDP. By default, no SDP is associated with a service. Once an SDP is created, services can be associated to that SDP.
- An SDP is not specific to any one service or any type of service. An SDP can have more than one service bound to it.
- The SDP IP address must be a 7750 SR system IP address.
- In order to configure an MPLS SDP, LSPs must be configured first and then the LSP-to-SDP association must be explicitly created.

To configure a basic SDP, perform the following steps:

1. Select an originating node.
2. Create an SDP ID.
1. Select an encapsulation type.
2. Select the far-end node.

To configure the return path SDP, perform the same steps on the far-end 7750 SR router.

1. Select an originating node.
2. Create an SDP ID.
3. Select an encapsulation type.
4. Select the far-end node.

Use the following CLI syntax to create an SDP and select an encapsulation type. If you do not specify GRE or MPLS, the default encapsulation type is GRE.

NOTE: When you specify the far-end ip address, you are creating the tunnel. In essence, you are creating the path from Point A to Point B. When you configure a distributed Epipe SAP, you must identify an SDP ID. Use the `show service sdp` command to display the qualifying SDPs.

CLI Syntax:
```
config>service# sdp sdp-id [gre | mpls] create
                description description-string
                far-end ip-addr
                lsp lsp-name [lsp-name]
                path-mtu octets
                no shutdown
                keep-alive
                        hello-time seconds
                        hold-down-time seconds
                        max-drop-count count
                        message-length octets
                        no shutdown
```

On the mirror-source router, configure an SDP pointing toward the mirror-destination router (or use an existing SDP):

Example:
```
*A:ALA-A>config>service# sdp 1 gre create
      config>service>sdp> description "to-GRE-10.10.10.104"
      config>service>sdp> far-end "10.10.10.104"
      config>service>sdp> no shutdown
```

On the mirror-destination router, configure an SDP pointing toward the mirror-source router (or use an existing SDP):

Example:
```
*A:ALA-B>config>service# sdp 4 gre create
      config>service>sdp> description "to-GRE-10.10.10.103
      config>service>sdp> far-end "10.10.10.103"
      config>service>sdp> no shutdown
```

The following example displays the SDP configurations on both the mirror-source and mirror-destination routers.

```
*A:ALA-A>config>service# info
----------------------------------------
       sdp 1 create
           description "to-10.10.10.104"
           far-end 10.10.10.104
           no shutdown
        exit
----------------------------------------
*A:ALA-A>config>service#

*A:ALA-B>config>service# info
----------------------------------------------
        sdp 4 create
            description "to-10.10.10.103"
            far-end 10.10.10.103
            no shutdown
        exit
----------------------------------------
*A:ALA-B>config>service#
```

Configuring a Remote Mirror Service

For remote mirroring, the source and destination are configured on the different routers. Note that mirror source and mirror destination parameters must be configured under the same service ID context.

The `mirror-source` commands are used as traffic selection criteria to identify traffic to be mirrored at the source. For example, use the `port` *`port-id`*`[.`*`channel-id`*`] {[egress] [ingress]}` and `mac-filter` *`mac-filter-id`* `entry` *`entry-id`* `[`*`entry-id`* `…]` commands.

Use the CLI syntax to configure one or more *mirror source* parameters:

CLI Syntax:
```
debug> mirror-source service-id
        ip-filter ip-filter-id entry entry-id [entry-id …]
        ingress-label label [label …]
        mac-filter mac-filter-id entry entry-id [entry-id …]
        port {port-id|lag lag-id} {[egress][ingress]}
        sap sap-id {[egress][ingress]}
        sdp sdp-id:[vc-id] {[egress] [ingress]}
        no shutdown
```

The `mirror-dest` commands are used to specify where the mirrored traffic is to be sent, the forwarding class, and the size of the packet. Use the following CLI syntax to configure mirror destination parameters:

CLI Syntax:
```
config>mirror#
        mirror-dest service-id [type {ether|frame-relay|ppp|atm-
        sdu}]
           description string
           fc fc-name
           remote-source
              far-end ip-addr ing-svc-label ing-svc-label
           sap sap-id
           sdp sdp-id[:vc-id][egr-svc-label [label|tldp]
           no shutdown
           slice-size bytes
```

The following displays the *mirror destination*, which is on ALA-A, configuration for mirror service 1216. This configuration specifies that the mirrored traffic coming from the *mirror source* (10.10.0.91) is to be directed to SAP 4/1/58 and states that the service only accepts traffic from far end 10.10.0.92 (ALA-B) with an ingress service label of 5678. When a forwarding class is specified, then all mirrored packets transmitted to the destination SAP or SDP override the default (be) forwarding class. The slice size limits the size of the stream of packet through the 7750 SR and the core network.

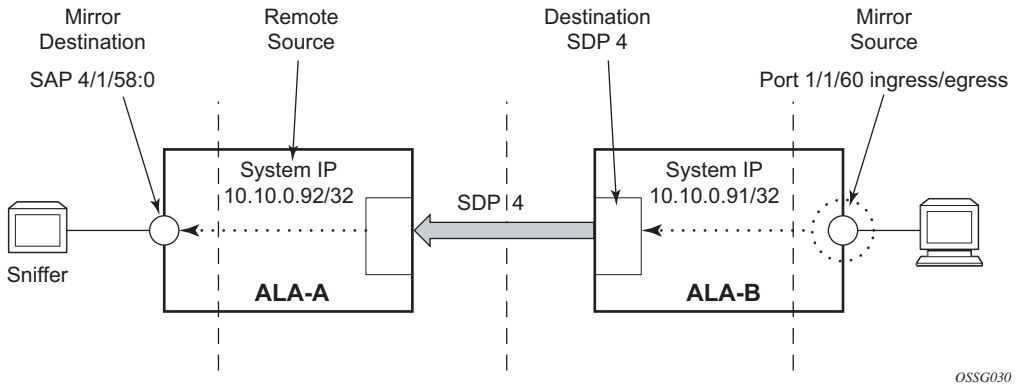

Figure 4: Remote Mirrored Service Tasks

The following example displays the CLI commands to configure remote mirrored service 1216. The traffic ingressing and egressing port 1/1/60 on 10.10.0.92 (ALA-B) will be mirrored to the destination SAP 4/1/58:0 on ALA-A.

Example:
```
*A:ALA-A>config>mirror# mirror-dest 1216 create
      config>mirror>mirror-dest$ description "Receiving mirrored
                                          traffic from .91"
      config>mirror>mirror-dest# sap 1/1/14:0 create
      config>mirror>mirror-dest>sap$ exit
      config>mirror>mirror-dest# no shutdown
      config>mirror>mirror-dest# remote-source
      config>mirror>mirror-dest>remote-source# far-end
                                          10.10.0.91 ing-svc-label 5678
      config>mirror>mirror-dest>remote-source# exit
```

Example: *A:ALA-B>config>mirror# mirror-dest 1216 create
    config>mirror>mirror-dest#description "Sending mirrored traffic to .92"
    config>mirror>mirror-dest# sdp 2 egr-svc-label 5678
    config>mirror>mirror-dest# fc h1
    config>mirror>mirror-dest# slice-size 128
    config>mirror>mirror-dest# exit all
    # debug mirror-source 1216
    debug>mirror-source# port 1/1/1.1.8 egress ingress
    debug>mirror-source# no shutdown

The following displays the *mirror destination* configuration for mirror service 1216 on ALA-A.

```
*A:ALA-A>config>mirror# info
----------------------------------------------
        mirror-dest 1216 create
            description "Receiving mirror traffic from .91"
            remote-source
                far-end 10.10.0.91 ing-svc-label 5678
            exit
            sap 1/1/14:0 create
                egress
                    qos 1
                exit
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-A>config>mirror#
```

The following displays the *remote mirror destination* configured on ALA-B:

```
*A:ALA-B>config>mirror# info
----------------------------------------------
        mirror-dest 1216 create
            description "Sending mirrored traffic to .92"
            fc h1
            sdp 2 egr-svc-label 5678
            slice-size 128
            no shutdown
        exit
----------------------------------------------
*A:ALA-B>config>mirror#
```

The following displays the *mirror source* configuration for ALA-B:

```
*A:ALA-B# show debug mirror
debug
    mirror-source 1216
        port 7/1/1.1.8 egress ingress
        no shutdown
```

```
    exit
exit
*A:ALA-B#
```

The following displays the SDP configuration from ALA-A to ALA-B (SDP 2) and the SDP configuration from ALA-B to ALA-A (SDP 4).

```
*A:ALA-A>config>service>sdp# info
----------------------------------------------
            description "GRE-10.10.0.91"
            far-end 10.10.0.01
            no shutdown
----------------------------------------------
*A:ALA-A>config>service>sdp#


*A:ALA-B>config>service>sdp# info
----------------------------------------------
            description "GRE-10.10.20.92"
            far-end 10.10.10.103
            no shutdown
----------------------------------------------
*A:ALA-B>config>service>sdp#
```

Configuring an ATM Mirror Service

Configure a local ATM mirror service at PE1:

Example:
```
config>mirror# mirror-dest 1 type atm-sdu create
    config>mirror>mirror-dest# sap 1/2/1:1/101 create
    config>mirror>mirror-dest>sap# no shutdown
    config>mirror>mirror-dest>sap# exit all
    # debug
    debug# mirror-source 1
    debug>mirror-source# sap 2/1/1/:0/100 ingress
```

Configure a remote ATM mirror service at PE1:

Example:
```
config>mirror# mirror-dest 1 type atm-sdu create
    config>mirror>mirror-dest# sdp 1:20
    config>mirror>mirror-dest# exit all
    # debug

    debug# mirror-source 1
    debug>mirror-source# sap 2/1/1/:0/100 ingress
```

Configure a remote ATM mirror service at PE2:

Example:
```
config>mirror# mirror-dest 1 type atm-sdu create
    config>mirror>mirror-dest# remote-source
    config>mirror>mirror-dest>remote-source# far-end 10.10.10.10
    config>mirror>mirror-dest>remote-source# exit
    config>mirror>mirror-dest# sap 1/2/1:1/101 create
```

Service Management Tasks

This section discusses the following service management tasks:

Use the following command syntax to modify an existing mirrored service:

CLI Syntax:
```
config>mirror#
 mirror-dest service-id [type {ether|frame-relay|ppp|atm-
 sdu}]
    description description-string
    no description
    fc fc-name
    no fc
    remote-source
       far-end ip-address [ing-svc-label ing-svc-label|tldp]
       no far-end ip-address
    sap sap-id
    no sap
    sdp sdp-name [egr-svc-label egr-svc-label|tldp]
    no sdp
    [no] shutdown
    slice-size bytes
    no slice-size
```

CLI Syntax:
```
debug
[no] mirror-source service-id
    ip-filter ip-filter-id entry entry-id [entry-id]
    no ip-filter ip-filter-id
    no ip-filter ip-filter-id entry entry-id [entry-id]
    ingress-label label [label]
    no ingress-label
    no ingress-label label [label]
    mac-filter mac-filter-id entry entry-id [entry-id]
    no mac-filter mac-filter-id
    no mac-filter mac-filter-id entry entry-id [entry-id]
    [no] port {port-id|lag lag-id} {[egress][ingress]}
    [no] sap sap-id {[egress] [ingress]}
    [no] shutdown
```

Modifying a Local Mirrored Service

Existing mirroring parameters can be modified in the CLI. The changes are applied immediately. The service must be shut down if changes to the SAP are made. Refer to page 1335 for the complete list of commands and options.

The following example displays commands to modify parameters for a basic local mirroring service.

Example:

```
config>mirror# mirror-dest 103
      config>mirror>mirror-dest# shutdown
      config>mirror>mirror-dest# no sap
      config>mirror>mirror-dest# sap 3/1/5:0 create
      config>mirror>mirror-dest>sap$ exit
      config>mirror>mirror-dest# fc premium
      config>mirror>mirror-dest# slice-size 128
      config>mirror>mirror-dest# no shutdown

      debug# mirror-dest 103
      debug>mirror-source# no port 2/1/24 ingress egress
      debug>mirror-source# port 3/1/7 ingress egress
```

The following displays the local mirrored service modifications:

```
*A:ALA-A>config>mirror# info
----------------------------------------------
mirror-dest 103 create
            no shutdown
            fc premium
            remote-source
            exit
            sap 3/1/5:0 create
                egress
                    qos 1
                exit
            exit
            slice-size 128
        exit

*A:ALA-A>debug>mirror-source# show debug mirror
debug
    mirror-source 103
        no shutdown
        port 3/1/7 egress ingress
    exit
*A:ALA-A>debug>mirror-source#
```

Deleting a Local Mirrored Service

Existing mirroring parameters can be deleted in the CLI. A shutdown must be issued on a service level in order to delete the service. It is not necessary to shut down or remove SAP or port references to delete a local mirrored service. Refer to page 1335 for the complete list of commands and options.

The following example displays commands to delete a local mirrored service.

Example:ALA-A>
```
config>mirror# mirror-dest 103
config>mirror>mirror-dest# shutdown
config>mirror>mirror-dest# exit
config>mirror# no mirror-dest 7
config>mirror# exit
```

Modifying a Remote Mirrored Service

Existing mirroring parameters can be modified in the CLI. The changes are applied immediately. The service must be shut down if changes to the SAP are made. Refer to page 1335 for the complete list of commands and options.

In the following example, the mirror destination is changed from 10.10.10.2 (ALA-B) to 10.10.10.3 (SR3). Note that the mirror-dest service ID on ALA-B must be shut down first before it can be deleted.

The following example displays commands to modify parameters for a remote mirrored service.

Example:

```
*A:ALA-A>config>mirror# mirror-dest 104
       config>mirror>mirror-dest# remote-source
       config>mirror>mirror-dest>remote-source# no far-end 10.10.10.2
       remote-source# far-end 10.10.10.3 ing-svc-label 3500

       *A:ALA-B>config>mirror# mirror-dest 104
       config>mirror>mirror-dest# shutdown
       config>mirror>mirror-dest# exit
       config>mirror# no mirror-dest 104

       SR3>config>mirror# mirror-dest 104 create
       config>mirror>mirror-dest# sdp 4 egr-svc-label 3500
       config>mirror>mirror-dest# no shutdown
       config>mirror>mirror-dest# exit all

       SR3>#  debug
       debug# mirror-source 104
       debug>mirror-source# port 5/1/2 ingress egress
       debug>mirror-source# no shutdown

*A:ALA-A>config>mirror# info
----------------------------------------------
       mirror-dest 104 create
           remote-source
               far-end 10.10.10.3 ing-svc-label 3500
           exit
           sap 2/1/15:0 create
               egress
                   qos 1
               exit
           exit
           no shutdown
       exit

SR3>config>mirror# info
----------------------------------------------
        mirror-dest 104 create
            sdp 4  egr-svc-label 3500
            no shutdown
        exit
```

```
-----------------------------------------------
SR3>config>mirror#

SR3# show debug mirror
debug
    mirror-source 104
        no shutdown
        port 5/1/2 egress ingress
    exit
exit
SR3#
```

Deleting a Remote Mirrored Service

Existing mirroring parameters can be deleted in the CLI. A shut down must be issued on a service level in order to delete the service. It is not necessary to shut down or remove SAP, SDP, or far-end references to delete a remote mirrored service. Refer to page 1335 for the complete list of commands and options.

Mirror destinations must be shut down first before they can be deleted.

Example:
```
*A:ALA-A>config>mirror# mirror-dest 105
      config>mirror>mirror-dest# shutdown
      config>mirror>mirror-dest# exit
      config>mirror# no mirror-dest 105
      config>mirror# exit

      *A:ALA-B>config>mirror# mirror-dest 105
      config>mirror>mirror-dest# shutdown
      config>mirror>mirror-dest# exit
      config>mirror# no mirror-dest 105
      config>mirror# exit
```

The mirror-destination service ID 105 was removed from the configuration on ALA-A and ALA-B, thus, does not appear in the `info` command output.

```
*A:ALA-A>config>mirror# info
----------------------------------------------

----------------------------------------------
*A:ALA-A>config>mirror# exit


*A:ALA-B>config>mirror# info
----------------------------------------------

----------------------------------------------
*A:ALA-B>config>mirror# exit
```

Since the mirror destination was removed from the configuration on ALA-B, the port information was automatically removed from the `debug mirror-source` configuration.

```
*A:ALA-B# show debug mirror
debug
exit
*A:ALA-B#
```

Mirror Service Command Reference

Command Hierarchies

Configuration Commands

```
config
    — mirror
        — mirror-dest service-id  [type {ether | frame-relay | ppp | atm-sdu}]
        — no mirror-dest service-id
            — description string
            — no description
            — fc fc-name
            — no fc
            — [no] remote-source
                — far-end ip-address [ing-svc-label ing-vc-label | tldp]
                — no far-end ip-address
            — sap sap-id
            — no sap
                — egress
                    — qos policy-id
                    — no qos
            — sdp sdp-id[:vc-id] [egr-svc-label label | tldp]
            — no sdp
            — slice-size bytes
            — no slice-size
            — [no] shutdown
```

Debug Commands

```
debug
    — [no] mirror-source mirror-dest-service-id
        — ingress-label label [label …up to 8 max]
        — no ingress-label [label [label …up to 8 max]]
        — ip-filter ip-filter-id entry entry-id [entry-id …]
        — no ip-filter  ip-filter-id [entry entry-id] [entry-id …]
        — mac-filter mac-filter-id entry entry-id [entry-id …]
        — no mac-filter mac-filter-id [entry entry-id...]
        — port {port-id | lag lag-id} {[egress] [ingress]}
        — no port {port-id | lag lag-id} [egress] [ingress]
        — sap sap-id {[egress] [ingress]}
        — no sap sap-id [egress] [ingress]
        — [no] shutdown
```

Show Commands

show
- — **debug mirror** [*application*]
- — **mirror mirror-dest** [*service-id*]
- — **service**
  - — **service-using mirror**

Configuration Commands

Generic Commands

description

Syntax **description** *string*
no description

Context config>mirror>mirror-dest

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

The **description** command is a text string to help you identify the content of the file.

The **no** form of the command removes the description string.

Default There is no default 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.

shutdown

Syntax [**no**] **shutdown**

Context config>mirror>mirror-dest
debug>mirror-source

Description The **shutdown** command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

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

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

The **no** form of the command puts an entity into the administratively enabled state.

Default See Special Cases below.

Special Cases **Mirror Destination —** When a mirror destination service ID is shutdown, mirrored packets associated with the service ID are not accepted from the mirror source or remote source 7750 SR7750 SR router. The associated mirror source is put into an operationally down mode. Mirrored packets are not transmitted out of the SAP or SDP. Each mirrored packet is silently discarded. If the mirror destination is a SAP, the SAP's discard counters are incremented.

The **shutdown** command places the mirror destination service or mirror source into an administratively down state. The **mirror-dest** service ID must be shut down in order to delete the service ID, SAP or SDP association from the system.

The default state for a mirror destination service ID is **shutdown**. A **no shutdown** command is required to enable the service.

Mirror Source — Mirror sources do not need to be shutdown in order to remove them from the system.

When a mirror source is **shutdown**, mirroring is terminated for all sources defined locally for the **mirror-dest** service ID. If the **remote-source** command has been executed on the **mirror-dest** associated with the shutdown **mirror-source**, mirroring continues for remote sources.

The default state for a mirror source for a given **mirror-dest** service ID is **no shutdown**. A **shutdown** command is required to disable mirroring from that mirror-source.

Mirror Destination Configuration

far-end

Syntax **far-end** *ip-address* [**ing-svc-label** *ing-vc-label* | **tldp**]
no far-end *ip-addr*

Context config>mirror>mirror-dest>remote-source

Description This command defines the remote device and configures parameters for mirror destination services on other devices allowed to mirror to the mirror destination service ID.

The **far-end** command is used within the context of the **remote-source** node. It allows the definition of accepted remote sources for mirrored packets to this *mirror-dest-service-id*. Up to 50 **far-end** sources can be specified. If a far end 7750 SR router has not been specified, packets sent to the router are discarded.

The **far-end** command is used to define a remote source 7750 SR that may send mirrored packets to this 7750 SR for handling by this **mirror-dest** *service-id*.

The **ing-svc-label** keyword must be given to manually define the expected ingress service label. This ingress label must also be manually defined on the far end address through the **mirror-dest** SDP binding keyword **egr-svc-label**.

The **no** form of the command deletes a far end address from the allowed remote senders to this **mirror-dest** service. All **far-end** addresses are removed when **no remote-source** is executed. All signaled ingress service labels are withdrawn from the far end address affected. All manually defined *ing-svc-label* are removed.

Default No far end service ingress addresses are defined.

Parameters *ip-address* — The service IP address (system IP address) of the remote 7750 SR device sending mirrored traffic to this mirror destination service. If 0.0.0.0 is specified, any remote 7750 SR is allowed to send to this service.

Values 1.0.0.1 — 223.255.255.254

The ingress service label must be manually defined using the **ing-svc-label** keyword. On the far end 7750 SR, the associated SDP **egr-svc-label** must be manually set and equal to the label defined in **ing-svc-label**.

ing-svc-label *ing-vc-label* **—** Specifies the ingress service label for mirrored service traffic on the **far end** device for manually configured mirror service labels.

The defined *ing-svc-label* is entered into the ingress service label table which causes ingress packet with that service label to be handled by this **mirror-dest** service.

The specified *ing-svc-label* must not have been used for any other service ID and must match the far end expected specific *egr-svc-label* for this 7750 SR. It must be within the range specified for manually configured service labels defined on this 7750 SR. It may be reused for other far end addresses on this *mirror-dest-service-id*.

Values 2047 — 18431

tldp — Specifies that the label is obtained through signaling via the LDP.

fc

Syntax **fc** *fc-name*
no fc

Context config>mirror>mirror-dest

Description This command specifies a forwarding class for all mirrored packets transmitted to the destination SAP or SDP overriding the default (be) forwarding class. All packets are sent with the same class of service to minimize out of sequence issues. The mirrored packet does not inherit the forwarding class of the original packet.

When the destination is on a SAP, a single egress queue is created that pulls buffers from the buffer pool associated with the *fc-name*.

When the destination is on an SDP, the *fc-name* defines the DiffServ based egress queue that will be used to reach the destination. The *fc-name* also defines the encoded forwarding class of the encapsulation.

The **no** form of the command reverts the **mirror-dest** service ID forwarding class to the default forwarding class.

Default The best effort (be) forwarding class is associated with the **mirror-dest** service ID.

Parameters *fc-name* — The name of the forwarding class with which to associate mirrored service traffic. The forwarding class name must already be defined within the system. If the fc-name does not exist, an error will be returned and the **fc** command will have no effect. If the *fc-name* does exist, the forwarding class associated with *fc-name* will override the default forwarding class.

Values `be`, `l2`, `af`, `l1`, `h2`, `ef`, `h1`, `nc`

mirror-dest

Syntax **mirror-dest** *service-id* [**type** {**ether** | **frame-relay** | **ppp | atm-sdu**}]
no mirror-dest

Context config>mirror

Description This command creates a context to set up a service that is intended for packet mirroring. It is configured as a service to allow mirrored packets to be directed locally (within the same 7750 SR router) or remotely, over the core of the network and have a far end 7750 SR decode the mirror encapsulation.

The **mirror-dest** service is comprised of destination parameters that define where the mirrored packets are to be sent. It also specifies whether the defined *service-id* will receive mirrored packets from far end 7750 SR devices over the network core.

Only 31 **mirror-dest** service IDs can be created without the **remote-only** flag specified within a single system.

The **mirror-dest** service IDs are persistent between boots of the router and are included in the configuration saves. The local sources of mirrored packets for the service ID are defined within the **debug mirror mirror-source** command that references the same *service-id*. Up to 255 **mirror-dest** service IDs can be created within a single system.

The **mirror-dest** command is used to create or edit a service ID for mirroring purposes. If the *service-id* does not exist within the context of all defined services, the **mirror-dest** service is created and the context of the CLI is changed to that service ID. If the *service-id* exists within the context of defined **mirror-dest** services, the CLI context is changed for editing parameters on that service ID. If the *service-id* exists within the context of another service type, an error message is returned and CLI context is not changed from the current context.

The actual packet mirror enabling commands defined in the **debug mirror mirror-source** commands are not preserved in configuration saves.

The **no** form of the command removes a mirror destination from the system. The **mirror-source** associations with the **mirror-dest** *service-id* do not need to be removed or shutdown first. The **mirror-dest** *service-id* must be shutdown before the service ID can be removed. When the service ID is removed, all **mirror-source** commands that have the service ID defined will also be removed from the system.

Default No packet mirroring services are defined.

Parameters *service-id* — The service identification identifies the service in the service domain. This ID is unique to this service and cannot be used by any other service, regardless of service type. The same service ID must be configured on every 7750 SR router that this particular service is defined on.

If particular a service ID already exists for a service, then the same value cannot be used to create a mirror destination service ID with the same value.

For example:

If an Epipe service-ID **11** exists, then a mirror destination service-ID **11** cannot be created. If a VPLS service-ID **12** exists, then a mirror destination service-ID **12** cannot be created.
If an IES service-ID **13** exists, then a mirror destination service-ID **13** cannot be created.

Values 1 — 2147483647

type — The type describes the encapsulation supported by the mirror service.

Values ether, frame-relay, ppp, atm-sdu

remote-source

Syntax [**no**] **remote-source**

Context config>mirror>mirror-dest

Description This command configures remote devices to mirror traffic to this device for mirror service egress. Optionally, deletes all previously defined remote mirror ingress devices.

The remote-source context allows the creation of a ‘sniffer farm’ to consolidate expensive packet capture and diagnostic tools to a central location. Remote areas of the access network can be monitored via normal service provisioning techniques.

Specific far-end 7750 SR devices can be specified with the **far-end** command allowing them to use this router as the destination for the same *mirror-dest-service-id*.

The **remote-source** node allows the source of mirrored packets to be on remote 7750 SR devices. The local 7750 SR will configure its network ports to forward packets associated with the *service-id*

to the destination SAP. When **remote-source far-end** addresses are configured, an SDP is not allowed as a destination.

By default, the **remote-source** context contains no **far-end** addresses. When no **far-end** addresses have been specified, network remote 7750 SR devices will not be allowed to mirror packets to the local 7750 SR as a mirror destination. Packets received from unspecified **far-end** addresses will be discarded at network ingress.

The **no** form of the command restores the *service-id* to the default condition to not allow a remote 7750 SR access to the mirror destination. The **far-end** addresses are removed without warning.

Default No remote source devices defined

sap

Syntax **sap** *sap-id*
no sap

Context config>mirror>mirror-dest

Description This command creates a service access point (SAP) within a mirror destination service. It also associates a predefined SAP within another service ID to a mirror source.

The SAP is defined with port and encapsulation parameters to uniquely identify the (mirror) SAP on the interface and within the box. The specified SAP must define a FastE, GigE, or XGigE access port with a dot1q, null, or q-in-q encapsulation type.

The mirror destination SAP referenced by the SAP ID. The SAP is owned by the mirror destination service ID. If the interface is administratively down, all SAPs on that interface are also operationally down. A SAP can only be defined on a port configured as an access port with the **mode** command at the interface level.

Only one SAP can be created within a **mirror-dest** service ID. If the defined SAP has not been created on any service within the system, the SAP is created and the context of the CLI will change to the newly created SAP.

If the defined SAP exists in the context of the service ID of the **mirror-dest** service, the CLI context is changed to the predefined SAP.

If the defined SAP exists in the context of another service ID, **mirror-dest** or any other type, an error is generated and the CLI context is not changed from the current context.

Mirror destination SAPs can be created on Ethernet interfaces that have been defined as an access interface. If the interface is defined as network, the SAP creation returns an error and the current CLI context is not changed.

When the **no** form of this command is used on a SAP created by a mirror destination service ID, the SAP with the specified port and encapsulation parameters is deleted. The SAP may not be deleted until the **mirror-dest** service ID been **shutdown**.

Default No default SAP for the mirror destination service defined.

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	*slot*/*mda*/*port*[.*channel*]	6/2/3.1
null	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]	*port-id*: 6/2/3 *bundle-id*: bundle-ima-10/1.1 bundle-ppp-10/1.1 *lag-id*: lag-100 *aps-id*: aps-1
dot1q	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]:qtag1	*port-id*:qtag1: 6/2/3:100 *bundle-id*:qtag1: bundle-ima-10/1.1:16 *bundle-id*:qtag1: bundle-ppp-10/1.1:16 *lag-id*:qtag1: lag-100:102 *aps-id*:qtag1: aps-1:103
qinq	[*port-id* | *bundle-id* | *lag-id*]:*qtag1.qtag2*	*port-id*:qtag1.qtag2: 6/2/3:100.10 *bundle-id*:qtag1.qtag2: bundle-ima-10/1.1:16.17 bundle-ppp-10/1.1:16.17 *lag-id*:qtag1.qtag2: lag-100: *bundle-id*:qtag1.qtag2: bundle-ima-5/1.1:100.10 bundle-ppp-5/1.1:100.10
frame-relay	[*port-id* | *aps-id*]:*dlci*	*port-id*: 9/1/1:100 *aps-id*: aps-1 *dlci*: 16
ima-grp	*bundle-id*[:*vpi/vci* | *vpi* | *vpi1.vpi2*]	*bundle-type-slot/mda.bundle-num*: bundle-ima-10/1.1:16 bundle-ppp-10/1.1:16 *vpi/vci*: 16/60 *vpi*: 16 *vpi1.vpi2*: 16.200

Values sap-id

null	port-id
dot1q	port-id:qtag1
qinq	port-id:qtag1.qtag2
atm	port-id:vpi/vci
frame	port-id

port-id	slot/mda/port[.channel]
qtag1	0 — 4094
qtag2	*, 0 — 4094
vpi	0 — 4095 (NNI) 0 — 255 (UNI)
vci	1, 2, 5 — 65535

Values

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 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

qtag1, qtag2 — Specifies the encapsulation value used to identify the SAP on the port or sub-port. If this parameter is not specificially defined, the default value is 0.

Values qtag1: 0 — 4094
qtag2 : * | 0 — 4094

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 TDM	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 TDM	BCP-Dot1q	0 — 4094	The SAP is identified by the 802.1Q tag on the channel.
SONET/SDH TDM	Frame Relay	16 — 991	The SAP is identified by the data link connection identifier (DLCI).

egress

Syntax **egress**

Context config>mirror>mirror-dest>sap

Description This command enables access to the context to associate an egress SAP Quality of Service (QoS) policy with a mirror destination SAP.

If no QoS policy is defined, the system default SAP egress QoS policy is used for egress processing.

qos

Syntax **qos** *policy-id*
no qos

Context config>mirror>mirror-dest>sap>egress

Description This command associates a QoS policy with an egress SAP for a mirrored service.

By default, no specific QoS policy is associated with the SAP for egress, so the default QoS policy is used.

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

Default QoS policy-id 1.

Parameters *policy-id* — The QoS policy ID to associate with SAP for the mirrored service. The policy ID must already exist.

Values 1 — 65535

sdp

Syntax **sdp** *sdp-id*[:*vc-id*] [**egr-svc-label** *label* **| tlpd**]
no sdp

Context config>mirror>mirror-dest

Description This command binds an existing (mirror) service distribution path (SDP) to the mirror destination service ID.

The operational state of the SDP dictates the operational state of the SDP binding to the mirror destination. If the SDP is shutdown or operationally down, then SDP binding is down. Once the binding is defined and the service and SDP are operational, the far-end 7750 SR defined in the **config service sdp** *sdp-id* **far-end** parameter is considered part of the service ID.

Only one SDP can be associated with a mirror destination service ID. If a second **sdp** command is executed after a successful SDP binding, an error occurs and the command has no effect on the existing configuration. A **no sdp** command must be issued before a new SDP binding can be attempted.

An SDP is a logical mechanism that ties a far end 7750 SR to a specific service without having to define the far-end SAP. Each SDP represents a method to reach a 7750 SR.

One method is the IP Generic Router Encapsulation (GRE) encapsulation, which has no state in the core of the network. GRE does not specify a specific path to a 7750 SR router. A GRE-based SDP uses the underlying IGP routing table to find the best next hop to the far end router.

The other method is Multi-Protocol Label Switching (MPLS) encapsulation. 7750 SR routers support both signaled and non-signaled LSPs (Label Switched Path) though the network. Non-signaled paths are defined at each hop through the network. Signaled paths are protocol communicated from end to end using RSVP. Paths may be manually defined or a constraint based routing protocol (i.e., OSPF-TE or CSPF) can be used to determine the best path with specific constraints.

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

An egress service label (Martini VC-Label), used by the SDP to differentiate each service bound to the SDP to the far-end router, must be obtained manually or though signaling with the far end. If manually configured, it must match the **ing-svc-label** defined for the local router.

The **no** form of the command removes the SDP binding from the mirror destination service. Once removed, no packets are forwarded to the far-end (destination) router from that mirror destination service ID.

Default No default SDP ID is bound to a mirror destination service ID. If no SDP is bound to the service, the mirror destination will be local and cannot be to another 7750 SR over the core network.

Parameters *sdp-id*[:*vc-id*] — A locally unique SDP identification (ID) number. The SDP ID must exist. If the sdp-id does not exist, an error will occur and the command will not execute.

For mirror services, the *vc-id* defaults to the *service-id*. However, there are scenarios where the *vc-id* is being used by another service. In this case, the SDP binding cannot be created. So, to avoid this, the mirror service SDP bindings now accepts *vc-ids*.

Values 1 — 17407

egr-svc-label *label* — The **egr-svc-label** keyword is used to define the *egr-svc-label* used to identify this **mirror-dest** over this *sdp-id*. The *egr-svc-label* must be explicitly configured.

The specified *egr-svc-label* must be locally unique within this 7750 SR and match the far end expected specific *ing-svc-label* for this 7750 SR. It must be within the range specified for manually configured service labels defined on this 7750 SR.

Default None — Must be explicitly configured.

Values 16 — 1048575

tldp — Specifies that the label is obtained through signaling via the LDP.

slice-size

Syntax **slice-size** *bytes*
no slice-size

Context config>mirror>mirror-dest

Description This command enables mirrored frame truncation and specifies the maximum size, in bytes, of a mirrored frame that can be transmitted to the mirror destination.

This command enables mirroring larger frames than the destination packet decode equipment can handle. It also allows conservation of mirroring resources by limiting the size of the packet stream through the router and the core network.

When defined, the mirror **slice-size** creates a threshold that truncates a mirrored frame to a specific size. For example, if the value of 256 bytes is defined, a frame larger than 256 bytes will only have the first 256 bytes transmitted to the mirror destination. The original frame is not affected by the truncation. The mirrored frame size may increase if encapsulation information is added during transmission through the network core or out the mirror destination SAP to the packet/protocol decode equipment.

The actual capability of the 7750 SR to transmit a sliced or non-sliced frame is also dictated by the mirror destination SDP **path-mtu** and/or the mirror destination SAP physical MTU. Packets that require a larger MTU than the mirroring destination supports are discarded if the defined **slice-size** does not truncate the packet to an acceptable size.

The **no** form of the command disables mirrored packet truncation.

Default **no slice-size** — Mirrored packet truncation is disabled.

Parameters *bytes* — The number of bytes to which mirrored frames will be truncated, expressed as a decimal integer.

Values 128 — 9216

Mirror Source Configuration

ingress-label

Syntax [**no**] **ingress-label** *label* [*label …*up to 8 max]
no ingress-label *label* [*label …*up to 8 max]

Context debug>mirror-source

Description This command enables ingress MPLS frame mirroring based on the top-of-stack MPLS label. Multiple labels can be defined simultaneously.

The **ingress-label** command is used to mirror ingressing MPLS frames with specific MPLS labels to a specific mirror destination. The ingress label must be at the top of the label stack and can only be mirrored to a single mirror destination. If the same label is defined with multiple mirror destinations, an error is generated and the original mirror destination remains.

The **ingress-label** mirror source overrides all other mirror source definitions. The MPLS frame is mirrored to the mirror destination as it is received on the ingress network port. The 7750 SR MPLS label space is global for the system. A specific label is mirrored to the mirror destination regardless of the ingress interface.

By default, no ingress MPLS frames are mirrored. The **ingress-label** command must be executed to start mirroring on a specific MPLS label.

The **no ingress-label** command removes all label mirroring for the mirror source. To stop mirroring on specific labels, use the **no ingress-label** *label* form of the command. Multiple labels may be given in a single **no ingress-label** command.

Default No ingress MPLS labels for mirroring are defined.

Parameters *label* — The top-of-stack label received on ingress to be mirrored. A label can only be mirrored to a single mirror destination.

If the label does not exist on any ingress network ports, no packets are mirrored for that label. An error will not occur. Once the label exists on a network port, ingress mirroring commences for that label.

Values 0 — 1048575. The local MPLS stack may not support portions of this range.

ip-filter

Syntax **ip-filter** *ip-filter-id* **entry** *entry-id* [*entry-id …*]
no ip-filter *ip-filter-id*
no ip-filter *ip-filter-id* **entry** *entry-id* [*entry-id …*]

Context debug>mirror-source

Description This command enables mirroring of packets that match specific entries in an existing IP filter.

The IP filter must be defined with an **exclusive** scope. This prevents mirroring issues when an IP filter template is defined on multiple interfaces. The **mac-filter** command defines mirroring for MAC based filters.

The **ip-filter** command directs packets which match the defined list of entry IDs to be mirrored to the mirror destination referenced by the *mirror-dest-service-id* of the **mirror-source**.

The IP filter must already exist in order for the command to execute. Filters are configured in the **config>filter** context. If the IP filter does not exist, an error will occur. If the filter exists but has not been associated with a SAP or IP interface, an error is not generated but mirroring will not be enabled (there are no packets to mirror). Once the IP filter is defined to a SAP or IP interface, mirroring is enabled.

If the IP filter is defined as ingress, only ingress packets are mirrored. Ingress mirrored packets are mirrored to the mirror destination prior to any ingress packet modifications.

If the IP filter is defined as egress, only egress packets are mirrored. Egress mirrored packets are mirrored to the mirror destination after all egress packet modifications.

An *entry-id* within an IP filter can only be mirrored to a single mirror destination. If the same *entry-id* is defined multiple times, an error occurs and only the first **mirror-source** definition is in effect.

By default, no packets matching any IP filters are mirrored. Mirroring of IP filter entries must be explicitly defined.

The **no ip-filter** command, without the **entry** keyword, removes mirroring on all *entry-id*'s within the *ip-filter-id*.

When the **no** command is executed with the **entry** keyword and one or more *entry-id*'s, mirroring of that list of *entry-id*'s is terminated within the *ip-filter-id*. If an *entry-id* is listed that does not exist, an error will occur and the command will not execute. If an *entry-id* is listed that is not currently being mirrored, no error will occur for that *entry-id* and the command will execute normally.

Default IP filter mirroring is not defined.

Parameters *ip-filter-id* — The IP filter ID whose entries are mirrored. The IP filter must have a defined scope of **exclusive**. If the *ip-filter-id* does not exist, an error will occur and the command will not execute. Mirroring of packets will commence once the *ip-filter-id* is defined on a SAP or IP interface.

entry *entry-id* [*entry-id* …] **—** The IP filter entries to use as match criteria for packet mirroring. The **entry** keyword begins a list of *entry-id*'s for mirroring. Multiple *entry-id* entries may be specified with a single command. Each *entry-id* must be separated by a space. Up to <N><n> entry IDs may be specified in a single command.

If an *entry-id* does not exist within the IP filter, an error occurs and the command will not execute.

If the filter's *entry-id* is renumbered within the IP filter definition, the old *entry-id* is removed but the new *entry-id* must be manually added to the configuration to include the new (renumbered) entry's criteria.

If no *entry-id* entries are specified in the command, mirroring will not occur for that IP filter ID. The command will have no effect.

mac-filter

Syntax **mac-filter** *mac-filter-id* **entry** *entry-id* [*entry-id …*]
no mac-filter *mac-filter-id*
no mac-filter *mac-filter-id* **entry** *entry-id* [*entry-id …*]

Context debug>mirror-source

Description This command enables mirroring of packets that match specific entries in an existing MAC filter.

The IP filter must be defined with an **exclusive** scope. This prevents mirroring issues when an IP filter template is defined on multiple interfaces. The **ip-filter** command defines mirroring for MAC based filters.

The **mac-filter** command directs packets which match the defined list of entry IDs to be mirrored to the mirror destination referenced by the *mirror-dest-service-id* of the **mirror-source**.

The MAC filter must already exist in order for the command to execute. Filters are configured in the config>filter context. If the MAC filter does not exist, an error will occur. If the filter exists but has not been associated with a SAP or IP interface, an error is not be generated but mirroring will not be enabled (there are no packets to mirror). Once the filter is defined to a SAP or MAC interface, mirroring is enabled.

If the MAC filter is defined as ingress, only ingress packets are mirrored. Ingress mirrored packets are mirrored to the mirror destination prior to any ingress packet modifications.

If the MAC filter is defined as egress, only egress packets are mirrored. Egress mirrored packets are mirrored to the mirror destination after all egress packet modifications.

An *entry-id* within a MAC filter can only be mirrored to a single mirror destination. If the same *entry-id* is defined multiple times, an error occurs and only the first **mirror-source** definition is in effect.

By default, no packets matching any MAC filters are mirrored. Mirroring of MAC filter entries must be explicitly defined.

The **no mac-filter** command, without the **entry** keyword, removes mirroring on all *entry-id*'s within the *mac-filter-id*.

When the **no** command is executed with the **entry** keyword and one or more *entry-id*'s, mirroring of that list of *entry-id*'s is terminated within the *mac-filter-id*. If an *entry-id* is listed that does not exist, an error will occur and the command will not execute. If an *entry-id* is listed that is not currently being mirrored, no error will occur for that *entry-id* and the command will execute normally.

Default No MAC filter mirroring defined.

Parameters *mac-filter-id* — The MAC filter ID whose entries are mirrored. The MAC filter must have a defined scope of **exclusive**. If the *mac-filter-id* does not exist, an error will occur and the command will not execute. Mirroring of packets will commence once the *mac-filter-id* is defined on a SAP.

entry *entry-id* [*entry-id …*] **—** The MAC filter entries to use as match criteria for packet mirroring. The **entry** keyword begins a list of *entry-id*'s for mirroring. Multiple *entry-id* entries may be specified with a single command. Each *entry-id* must be separated by a space. Up to <N><n> entry IDs may be specified in a single command.

If an *entry-id* does not exist within the MAC filter ID, an error will occur and the whole command will not execute.

Each *entry-id* must exist within the *mac-filter-id*. If the *entry-id* is renumbered within the MAC filter definition, the old *entry-id* is removed from the list and the new *entry-id* will need to be manually added to the list if mirroring is still desired.

If no *entry-id* entries are specified in the command, mirroring will not occur for that MAC filter ID. The command will have no effect.

mirror-source

Syntax [**no**] **mirror-source** *service-id*

Context debug

Description This command configures mirror source parameters for a mirrored service.

The **mirror-source** command is used to enable mirroring of packets specified by the association of the **mirror-source** to sources of packets defined within the context of the *mirror-dest-service-id*. The mirror destination service must already exist within the system.

A mirrored packet cannot be mirrored to multiple destinations. If a mirrored packet is properly referenced by multiple mirror sources (for example, a SAP on one **mirror-source** and a port on another **mirror-source**), then the packet is mirrored to a single *mirror-dest-service-id* based on the following hierarchy:

1. Filter entry
2. MPLS label
3. Service access port (SAP)
4. Physical port

The hierarchy is structured so the most specific match criteria has precedence over a less specific match. For example, if a **mirror-source** defines a port and a SAP on that port, then the SAP mirror-source is accepted and the mirror-source for the port is ignored because of the hierarchical order of precedence.

The **mirror-source** configuration is not saved when a configuration is saved. A **mirror-source** manually configured within an ASCII configuration file will not be preserved if that file is overwritten by a **save** command. Define the **mirror-source** within a file associated with a **config exec** command to make a **mirror-source** persistent between system reboots.

By default, all **mirror-dest** service IDs have a **mirror-source** associated with them. The **mirror-source** is not technically created with this command. Instead the service ID provides a contextual node for storing the current mirroring sources for the associated **mirror-dest** service ID. The **mirror-source** is automatically created when the **mirror-dest** service ID is created. The **mirror-source** is also automatically removed when the **mirror-dest** service ID is deleted from the system.

The **no** form of the command deletes all related source commands within the context of the **mirror-source** *service-id*. The command does not remove the service ID from the system.

Default No mirror source match criteria is defined for the mirror destination service.

Parameters *service-id* — The mirror destination service ID for which match criteria will be defined. The *service-id* must already exist within the system.

Values 1 — 2147483647. The service ID must already exist as a **mirror-dest**.

port

Syntax **port** {*port-id* | **lag** *lag-id*} {[**egress**] [**ingress**]}
no port {*port-id* | **lag** *lag-id*} [**egress**] [**ingress**]

Context debug>mirror-source

Description This command enables mirroring of traffic ingressing or egressing a port (Ethernet port, SONET/SDH channel, TDM channel, or Link Aggregation Group (LAG)).

The **port** command associates a port or LAG to a mirror source. The port is identified by the *port-id*. The defined port may be Ethernet, SONET/SDH, access or network, or TDM channel, access. A network port may be a single port or a Link Aggregation Group (LAG) ID. When a LAG ID is given as the *port-id*, mirroring is enabled on all ports making up the LAG. If the port is a SONET/SDH interface, the *channel-id* must be specified to identify which channel is being mirrored. Either a LAG port member *or* the LAG port can be mirrored.

The port is only referenced in the mirror source for mirroring purposes. The mirror source association does not need to be removed before deleting the card to which the the port belongs. If the port is removed from the system, the mirroring association will be removed from the mirror source.

The same port may not be associated with multiple mirror source definitions with the **ingress** parameter defined. The same port may not be associated with multiple mirror source definitions with the **egress** parameter defined.

If a SAP is mirrored on an access port, the SAP mirroring will have precedence over the access port mirroring when a packet matches the SAP mirroring criteria. Filter and label mirroring destinations will also precedence over a port-mirroring destination.

If the port is not associated with a **mirror-source**, packets on that port will not be mirrored. Mirroring may still be defined for a SAP, label or filter entry, which will mirror based on a more specific criteria.

The encapsulation type on an access port or channel cannot be changed to Frame Relay if it is being mirrored.

The **no port** command disables port mirroring for the specified port. Mirroring of packets on the port may continue due to more specific mirror criteria. If the **egress** or **ingress** parameter keywords are specified in the **no** command, only the ingress or egress mirroring condition will be removed.

Default No ports are defined.

Parameters *port-id* — Specifies the port ID.

Syntax: port-id *slot/mda/port*[.*channel*]
aps-id aps-*group-id*[.*channel*]
aps keyword
group-id 1 — 64
bundle-id bundle-*slot/mda.bundle-num*
bundle keyword
bundle-num 1— 56
ccag-id ccag-*id.path-id*[*cc-type*]:*cc-id*
ccag keyword
id 1 — 8
path-id a, b
cc-type .sap-net, .net-sap
cc-id 0 — 4094

lag-id — The LAG identifier, expressed as a decimal integer.

egress — Specifies that packets egressing the port should be mirrored. Egress packets are mirrored to the mirror destination after egress packet modification.

ingress — Specifies that packets ingressing the port should be mirrored. Ingress packets are mirrored to the mirror destination prior to ingress packet modification.

sap

Syntax **sap** *sap-id* {[**egress**] [**ingress**]}
no sap *sap-id* [**egress**] [**ingress**]

Context debug>mirror-source

Description This command enables mirroring of traffic ingressing or egressing a service access port (SAP).

The mirror source SAP referenced by the *port-id* is owned by the service ID of the service in which it was created. The SAP is only referenced in the mirror source name for mirroring purposes. The mirror source association does not need to be removed before deleting the SAP from its service ID. If the SAP is deleted from its service ID, the mirror association is removed from the mirror source.

More than one SAP can be associated within a single **mirror-source**. Each SAP has its own **ingress** and **egress** parameter keywords to define which packets are mirrored to the mirror destination.

The SAP must be valid and properly configured. If the associated SAP does not exist, an error occurs and the command will not execute.

If the associated SAP has a native frame type (Ethernet or SONET/SDH) that is different from the **mirror-destination** *service-id* type, then the **disable-bcp-encap** parameter defined in the **destination** *service-id* will dictate whether frame encapsulation occurs.

The same SAP cannot be associated with multiple mirror source definitions for ingress packets. The same SAP may not be associated with multiple mirror source definitions for egress packets.

If a particular SAP is not associated with a mirror source name, then that SAP will not have mirroring enabled for that mirror source.

The **no** form of the command disables mirroring for the specified SAP. All mirroring for that SAP on ingress and egress is terminated. Mirroring of packets on the SAP can continue if more specific mirror criteria is configured. If the **egress** or **ingress** parameter keywords are specified in the **no** command, only the ingress or egress mirroring condition is removed.

Default No SAPs are defined by default.

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	*slot*/*mda*/*port*[.*channel*]	6/2/3.1
null	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]	*port-id*: 6/2/3 *bundle-id*: bundle-ima-10/1.1 bundle-ppp-10/1.1 *lag-id*: lag-100 *aps-id*: aps-1
dot1q	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]:qtag1	*port-id*:qtag1: 6/2/3:100 *bundle-id*:qtag1: bundle-ima-10/1.1:16 *bundle-id*:qtag1: bundle-ppp-10/1.1:16 *lag-id*:qtag1: lag-100:102 *aps-id*:qtag1: aps-1:103
qinq	[*port-id* | *bundle-id* | *lag-id*]:*qtag1.qtag2*	*port-id*:qtag1.qtag2: 6/2/3:100.10 *bundle-id*:qtag1.qtag2: bundle-ima-10/1.1:16.17 bundle-ppp-10/1.1:16.17 *lag-id*:qtag1.qtag2: lag-100: *bundle-id*:qtag1.qtag2: bundle-ima-5/1.1:100.10 bundle-ppp-5/1.1:100.10
atm	[*port-id* | *aps-id*][:*vpi/vci*|*vpi*|*vpi1.vpi2*]	*port-id*: 9/1/1:100/100 *aps-id*: aps-1 *vpi/vci*: 16/60 *vpi*: 16 *vpi1.vpi2*: 16.200
frame-relay	[*port-id* | *aps-id*]:*dlci*	*port-id*: 9/1/1:100 *aps-id*: aps-1 *dlci*: 16
cisco-hdlc	*slot/mda/port.channel*	*port-id*: 2/2/3.1
ima-grp	*bundle-id*[:*vpi/vci* | *vpi* | *vpi1.vpi2*]	*bundle-type-slot/mda.bundle-num*: bundle-ima-10/1.1:16 bundle-ppp-10/1.1:16 *vpi/vci*: 16/60 *vpi*: 16 *vpi1.vpi2*: 16.200

Values *sap-id*:

null	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *bundle-id*][:*vpi/vci* | *vpi* | *vpi1.vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*
ima-grp	*bundle-id*[:*vpi/vci* | *vpi* | *vpi1.vpi2*]

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 64
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

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 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

bundle-id — Specifies the multilink bundle to be associated with this IP interface. The **bundle** keyword must be entered at the beginning of the parameter.
The command syntax must be configured as follows:

bundle-id: **bundle**-*type-slot-id*/*mda-slot*.*bundle-num*
bundle-id value range: 1 — 128

For example:

```
*A:ALA-12>config#  port bundle-ima-5/1.1
*A:ALA-12>config>port# multilink-bundle
```

ima — Specifies Inverse Multiplexing over ATM. An IMA group is a collection of physical links bundled together and assigned to an ATM port.

qtag1, qtag2 — Specifies the encapsulation value used to identify the SAP on the port or sub-port. If this parameter is not specificially defined, the default value is 0.

Values qtag1: 0 — 4094
qtag2 : * | 0 — 4094

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 TDM	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 TDM	BCP-Dot1q	0 — 4094	The SAP is identified by the 802.1Q tag on the channel.
SONET/SDH TDM	Frame Relay	16 — 991	The SAP is identified by the data link connection identifier (DLCI).
SONET/SDH ATM	ATM	vpi (NNI) 0 — 4095 vpi (UNI) 0 — 255 vci 1, 2, 5 — 65535	The SAP is identified by the PVC identifier (vpi/vci).

channel-id — The SONET/SDH or TDM channel on the port of the SAP. A period separates the physical port from the *channel-id*. The port must be configured as an access port.

egress — Specifies that packets egressing the SAP should be mirrored. Egress packets are mirrored to the mirror destination after egress packet modification.

ingress — Specifies that packets ingressing the SAP should be mirrored. Ingress packets are mirrored to the mirror destination prior to ingress packet modification.

Show Commands

mirror

Syntax **mirror** [*application*]

Context **show>debug**

Description Displays mirror configuration and operation information.

Parameters *application* — Display specified protocol information.

ip, ospf, ospf3, rip, isis, mpls, rsvp, bgp, ldp, mirror, vrrp, frame-relay, igmp, pim, mtrace, system, filter, subscriber-mgmt, radius

service-using

Syntax **service-using** [**mirror**]

Context show>service

Description Displays mirror services.

If no optional parameters are specified, all services defined on the system are displayed.

Parameters **mirror** — Displays mirror services.

Output **Show Service-Using Mirror —** The following table describes service-using mirror 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-48# show service service-using mirror
===============================================================================
Services [mirror]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
```

```
218          Mirror    Up     Down       1                  04/20/2006 13:49:57
318          Mirror    Down   Down       1                  04/20/2006 13:49:57
319          Mirror    Up     Down       1                  04/20/2006 13:49:57
320          Mirror    Up     Down       1                  04/20/2006 13:49:57
1000         Mirror    Down   Down       1                  04/20/2006 13:49:57
1216         Mirror    Up     Down       1                  04/20/2006 13:49:57
1412412      Mirror    Down   Down       1                  04/20/2006 13:49:57
-------------------------------------------------------------------------------
Matching Services : 7
-------------------------------------------------------------------------------
===============================================================================
A:ALA-48#
```

mirror mirror-dest

Syntax **mirror mirror-dest *service-id***

Context show

Description Displays mirror configuration and operation information.

Parameters *service-id* — Specify the mirror service ID.

Output **Mirroring Output —** The following table describes the mirroring output fields:

Table 1: Mirroring Output Fields

Label	Description
Service Id	The service ID associated with this mirror destination.
Type	Entries in this table have an implied storage type of 'volatile'. The configured mirror source information is not persistent.
Admin State	Up — The mirror destination is administratively enabled.
	Down — The mirror destination is administratively disabled.
Oper State	Up — The mirror destination is operationally enabled.
	Down — The mirror destination is operationally disabled.
Forwarding Class	The forwarding class for all packets transmitted to the mirror destination.
Remote Sources	Yes — A remote source is configured.
	No — A remote source is not configured.
Slice	The value of the slice-size, i.e., the maximum portion of the mirrored frame that will be transmitted to the mirror destination. Any frame larger than the slice-size will be truncated to this value before transmission to the mirror destination. A value of 0 indicates that mirrored packet truncation based on slice size is disabled.

Table 1: Mirroring Output Fields (Continued)

Label	Description (Continued)
Destination SAP	The ID of the access port where the Service Access Point (SAP) associated with this mirror destination service is defined.
Egr QoS Policy	This value indicates the egress QoS policy ID. A value of 0 indicates that no QoS policy is specified.

Sample Output

```
A:SR7# show mirror mirror-dest 1000
===============================================================================
Mirror Service
===============================================================================
Service Id       : 1000                Type          : Ether
Admin State      : Up                  Oper State    : Down
Forwarding Class : be                  Remote Sources: No
Slice            : 0
Destination SAP  : 1/1/1               Egr QoS Policy: 1
-------------------------------------------------------------------------------
Local Sources
-------------------------------------------------------------------------------
Admin State      : Up
 - Port            1/1/2                          Egress Ingress
===============================================================================
A:SR7#

A:ALA-123>config>mirror# show mirror mirror-dest 500
===============================================================================
Mirror Service
===============================================================================
Service Id       : 500                 Type          : Ether
Admin State      : Up                  Oper State    : Up
Forwarding Class : be                  Remote Sources: Yes
Destination SAP  : 1/1/2               Egr QoS Policy: 1
-------------------------------------------------------------------------------
Remote Sources
-------------------------------------------------------------------------------
Far End          : 10.20.1.45          Ingress Label : 131070
-------------------------------------------------------------------------------
Local Sources
-------------------------------------------------------------------------------
Admin State      : Up
No Mirror Sources configured
===============================================================================
A:ALA-123>config>mirror#


A:ALA-456# show mirror mirror-dest 500
===============================================================================
Mirror Service
===============================================================================
Service Id       : 500                 Type          : Ether
Admin State      : Up                  Oper State    : Up
Forwarding Class : be                  Remote Sources: No
Destination SDP  : 144 (10.20.1.44)    Egress Label  : 131070
Signaling:       : TLDP
```

```
-------------------------------------------------------------------------------
Local Sources
-------------------------------------------------------------------------------
Admin State      : Up

No Mirror Sources configured
===============================================================================
A:ALA-456#
```

OAM and SAA

In This Chapter

This chapter provides information about the Operations, Administration and Management (OAM) and Service Assurance Agent (SAA) commands available in the CLI for troubleshooting services.

Topics in this chapter include:

OAM Overview

Delivery of services requires a number of operations occur properly and at different levels in the service delivery model. For example, operations such as the association of packets to a service, VC-labels to a service and each service to a service tunnel must be performed properly in the forwarding plane for the service to function properly. In order to verify that a service is operational, a set of in-band, packet-based Operation, Administration, and Maintenance (OAM) tools is required, with the ability to test each of the individual packet operations.

For in-band testing, the OAM packets closely resemble customer packets to effectively test the customer's forwarding path, but they are distinguishable from customer packets so they are kept within the service provider's network and not forwarded to the customer.

The suite of OAM diagnostics supplement the basic IP ping and traceroute operations with diagnostics specialized for the different levels in the service delivery model. There are diagnostics for MPLS LSPs, SDPs, Services and VPLS MACs within a service.

LSP Diagnostics

The 7750 SR OS LSP diagnostics are implementations of LSP ping and LSP traceroute based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. In an LDP ECMP network, a unique-path trace can be accomplished by specifying a unique 127/8 IP address for the **path-destination** *ip-address* parameter. Note that the 7750 SR can send multipath type 0 or 8, and up to a maximum of 36 bytes for multipath length (refer to RFC 4379 for more details). The 7750 SR supports unique-path trace on an LER of an LDP ECMP path. LSP ping, as described in the draft, provides a mechanism to detect dataplane failures in MPLS LSPs. LSP ping and LSP traceroute are modeled after the ICMP echo request/reply used by ping and traceroute to detect and localize faults in IP networks.

For a given FEC, LSP ping verifies whether the packet reaches the egress label edge router (LER), while in LSP traceroute mode, the packet is sent to the control plane of each transit label switched router (LSR) which performs various checks to see if it is actually a transit LSR for the path.

SDP Diagnostics

The 7750 SR OS SDP diagnostics are SDP ping and SDP MTU path discovery.

SDP Ping

SDP ping performs in-band uni-directional or round-trip connectivity tests on SDPs. The SDP ping OAM packets are sent in-band, in the tunnel encapsulation, so it will follow the same path as traffic within the service. The SDP ping response can be received out-of-band in the control plane, or in-band using the data plane for a round-trip test.

For a uni-directional test, SDP ping tests:

- Egress SDP ID encapsulation
- Ability to reach the far-end IP address of the SDP ID within the SDP encapsulation
- Path MTU to the far-end IP address over the SDP ID
- Forwarding class mapping between the near-end SDP ID encapsulation and the far-end tunnel termination

For a round-trip test, SDP ping uses a local egress SDP ID and an expected remote SDP ID. Since SDPs are uni-directional tunnels, the remote SDP ID must be specified and must exist as a configured SDP ID on the far-end 7750 SR. SDP round trip testing is an extension of SDP connectivity testing with the additional ability to test:

- Remote SDP ID encapsulation
- Potential service round trip time
- Round trip path MTU
- Round trip forwarding class mapping

SDP MTU Path Discovery

In a large network, network devices can support a variety of packet sizes that are transmitted across its interfaces. This capability is referred to as the Maximum Transmission Unit (MTU) of network interfaces. It is important to understand the MTU of the entire path end-to-end when provisioning services, especially for virtual leased line (VLL) services where the service must support the ability to transmit the largest customer packet.

The Path MTU Discovery tool provides a powerful tool that enables service provider to get the exact MTU supported between the service ingress and service termination points (accurate to one byte).

Service Diagnostics

Alcatel-Lucent's Service ping feature provides end-to-end connectivity testing for an individual service. Service ping operates at a higher level than the SDP diagnostics in that it verifies an individual service and not the collection of services carried within an SDP.

Service ping is initiated from a 7750 SR router to verify round-trip connectivity and delay to the far-end of the service. Alcatel-Lucent's implementation functions for both GRE and MPLS tunnels and tests the following from edge-to-edge:

- Tunnel connectivity
- VC label mapping verification
- Service existence
- Service provisioned parameter verification
- Round trip path verification
- Service dynamic configuration verification

VPLS MAC Diagnostics

While the LSP ping, SDP ping and Service ping tools enable transport tunnel testing and verify whether the correct transport tunnel is used, they do not provide the means to test the learning and forwarding functions on a per-VPLS-service basis.

It is conceivable, that while tunnels are operational and correctly bound to a service, an incorrect Forwarding Information Base (FIB) table for a service could cause connectivity issues in the service and not be detected by the ping tools. Alcatel-Lucent has developed VPLS OAM functionality to specifically test all the critical functions on a per-service basis. These tools are based primarily on the IETF document draft-stokes-vkompella-ppvpn-hvpls-oam-xx.txt, *Testing Hierarchical Virtual Private LAN Services*.

The VPLS OAM tools are:

- MAC Ping — Provides the ability to trace end-to-end switching of specified MAC addresses. MAC ping provides an end-to-end test to identify the egress customer-facing port where a customer MAC was learned. MAC ping can also be used with a broadcast MAC address to identify all egress points of a service for the specified broadcast MAC.
- MAC Trace — Provides the ability to trace a specified MAC address hop-by-hop until the last node in the service domain.
- CPE Ping — Provides the ability to check network connectivity to the specified client device within the VPLS. CPE ping will return the MAC address of the client, as well as the SAP and PE at which it was learned.

- MAC Populate — Allows specified MAC addresses to be injected in the VPLS service domain. This triggers learning of the injected MAC address by all participating nodes in the service. This tool is generally followed by MAC ping or MAC trace to verify if correct learning occurred.
- MAC Purge — Allows MAC addresses to be flushed from all nodes in a service domain.

MAC Ping

For a MAC ping test, the destination MAC address (unicast or multicast) to be tested must be specified. A MAC ping packet can be sent through the control plane or the data plane. When sent by the control plane, the ping packet goes directly to the destination IP in a UDP/IP OAM packet. If it is sent by the data plane, the ping packet goes out with the data plane format.

In the control plane, a MAC ping is forwarded along the flooding domain if no MAC address bindings exist. If MAC address bindings exist, then the packet is forwarded along those paths (if they are active). Finally, a response is generated only when there is an egress SAP binding to that MAC address. A control plane request is responded to via a control reply only.

In the data plane, a MAC ping is sent with a VC label TTL of 255. This packet traverses each hop using forwarding plane information for next hop, VC label, etc. The VC label is swapped at each service-aware hop, and the VC TTL is decremented. If the VC TTL is decremented to 0, the packet is passed up to the management plane for processing. If the packet reaches an egress node, and would be forwarded out a customer facing port, it is identified by the OAM label below the VC label and passed to the management plane.

MAC pings are flooded when they are unknown at an intermediate node. They are responded to only by the egress nodes that have mappings for that MAC address.

MAC Trace

A MAC trace functions like an LSP trace with some variations. Operations in a MAC trace are triggered when the VC TTL is decremented to 0.

Like a MAC ping, a MAC trace can be sent either by the control plane or the data plane.

For MAC trace requests sent by the control plane, the destination IP address is determined from the control plane mapping for the destination MAC. If the destination MAC is known to be at a specific remote site, then the far-end IP address of that SDP is used. If the destination MAC is not known, then the packet is sent unicast, to all SDPs in the service with the appropriate squelching.

A control plane MAC traceroute request is sent via UDP/IP. The destination UDP port is the LSP ping port. The source UDP port is whatever the system gives (note that this source UDP port is really the demultiplexor that identifies the particular instance that sent the request, when correlating the reply). The source IP address is the system IP of the sender.

When a traceroute request is sent via the data plane, the data plane format is used. The reply can be via the data plane or the control plane.

A data plane MAC traceroute request includes the tunnel encapsulation, the VC label, and the OAM, followed by an Ethernet DLC, a UDP and IP header. If the mapping for the MAC address is known at the sender, then the data plane request is sent down the known SDP with the appropriate tunnel encapsulation and VC label. If it is not known, then it is sent down every SDP (with the appropriate tunnel encapsulation per SDP and appropriate egress VC label per SDP binding).

The tunnel encapsulation TTL is set to 255. The VC label TTL is initially set to the min-ttl (default is 1). The OAM label TTL is set to 2. The destination IP address is the all-routers multicast address. The source IP address is the system IP of the sender.

The destination UDP port is the LSP ping port. The source UDP port is whatever the system gives (note that this source UDP port is really the demultiplexor that identifies the particular instance that sent the request, when correlating the reply).

The Reply Mode is either 3 (i.e., reply via the control plane) or 4 (i.e., reply through the data plane), depending on the reply-control option. By default, the data plane request is sent with Reply Mode 3 (control plane reply).

The Ethernet DLC header source MAC address is set to either the system MAC address (if no source MAC is specified) or to the specified source MAC. The destination MAC address is set to the specified destination MAC. The EtherType is set to IP.

CPE Ping

The MAC ping OAM tool makes it possible to detect whether a particular MAC address has been learned in a VPLS.

The **cpe-ping** command extends this capability to detecting end-station IP addresses inside a VPLS. A CPE ping for a specific destination IP address within a VPLS will be translated to a MAC-ping towards a broadcast MAC address. Upon receiving such a MAC ping, each peer PE within the VPLS context will trigger an ARP request for the specific IP address. The PE receiving a response to this ARP request will report back to the requesting 7750 SR.

MAC Populate

MAC Populate is used to send a message through the flooding domain to learn a MAC address as if a customer packet with that source MAC address had flooded the domain from that ingress point in the service. This allows the provider to craft a learning history and engineer packets in a particular way to test forwarding plane correctness.

The MAC populate request is sent with a VC TTL of 1, which means that it is received at the forwarding plane at the first hop and passed directly up to the management plane. The packet is

then responded to by populating the MAC address in the forwarding plane, like a conventional learn although the MAC will be an OAM-type MAC in the FIB to distinguish it from customer MAC addresses.

This packet is then taken by the control plane and flooded out the flooding domain (squelching appropriately, the sender and other paths that would be squelched in a typical flood).

This controlled population of the FIB is very important to manage the expected results of an OAM test.

The same functions are available by sending the OAM packet as a UDP/IP OAM packet. It is then forwarded to each hop and the management plane has to do the flooding.

Options for MAC Populate are to force the MAC in the table to type OAM (in case it already existed as dynamic or static or an OAM induced learning with some other binding), to prevent new dynamic learning to over-write the existing OAM MAC entry, to allow customer packets with this MAC to either ingress or egress the network, while still using the OAM MAC entry.

Finally, an option to flood the MAC Populate request causes each upstream node to learn the MAC (i.e., populate the local FIB with an OAM MAC entry), and to flood the request along the data plane using the flooding domain.

An age can be provided to age a particular OAM MAC after a different interval than other MACs in a FIB.

MAC Purge

MAC Purge is used to clear the FIBs of any learned information for a particular MAC address. This allows one to do a controlled OAM test without learning induced by customer packets. In addition to clearing the FIB of a particular MAC address, the purge can also indicate to the control plane not to allow further learning from customer packets. This allows the FIB to be clean, and be populated only via a MAC Populate.

MAC Purge follows the same flooding mechanism as the MAC Populate.

A UDP/IP version of this command is also available that does not follow the forwarding notion of the flooding domain, but the control plane notion of it.

VLL Diagnostics

VCCV Ping

VCCV ping is used to check connectivity of a VLL in-band. It checks that the destination (target) PE is the egress for the Layer 2 FEC. It provides a cross-check between the data plane and the control plane. It is in-band, meaning that the VCCV ping message is sent using the same encapsulation and along the same path as user packets in that VLL. This is equivalent to the LSP ping for a VLL service. VCCV ping reuses an LSP ping message format and can be used to test a VLL configured over an MPLS and GRE SDP.

VCCV-Ping Application

VCCV creates an IP control channel within the pseudowire between PE1 and PE2. PE2 should be able to distinguish, on the receive side, VCCV control messages from user packets on that VLL. The 7750 SR uses the router alert label immediately above the VC label to identify the VCCV-ping message. This method has a drawback that if ECMP is applied to the outer LSP label, such as the transport label, the VCCV message will not follow the same path as the user packets.

When sending the label mapping message for the VLL, PE1 and PE2 include an optional VCCV TLV in the PW FEC interface parameter field. The TLV indicates that the control channel will make use of the router alert label method. An example of VCCV ping is shown in Figure 71.

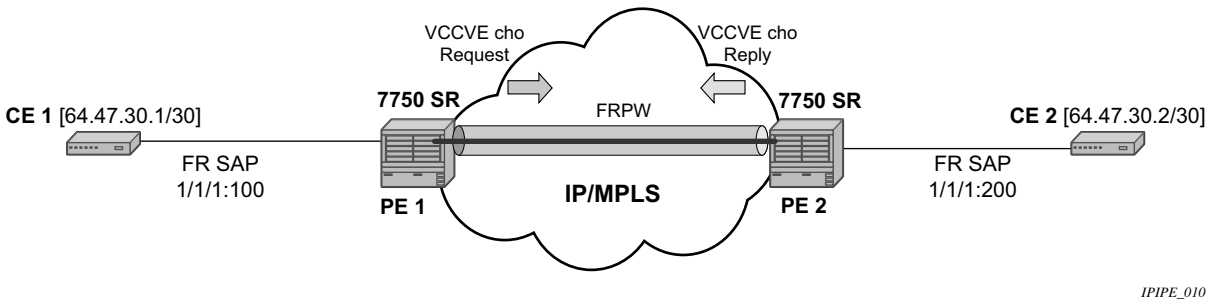

Figure 1: VCCV-Ping Application

A VCCV-ping is an LSP echo request message as defined in the LSP ping specification. It contains a Layer 2 FEC stack TLV in which it must include the sub-TLV type 10 FEC 128 pseudowire. It also contains a field which indicates to the destination PE which reply mode to use. There are four reply modes defined:

The 7750 SR supports the following reply modes:

- Reply by an IPv4/IPv6 UDP packet.
- Reply by application-level control channel. This mode sends the reply message in-band over the PW from PE2 to PE1. PE2 will encapsulate the echo reply message using the CC type negotiated with PE1. This is the default mode of operation.

The reply is an LSP echo reply message as defined in the specification. The message is sent as per the reply mode requested by PE1. The return codes supported are the same as those supported today in the 7750 SR LSP ping capability

The VCCV ping feature is in addition to the service ping OAM feature which can be used to test a service between 7750 SR nodes. The VCCV ping feature can test connectivity of a VLL with any third party node which is compliant to draft-ietf-pwe3-vccv-xx.txt.

VCCV-Ping in a Multi-Segment Pseudowire

Figure 72 displays and example of an application of VCCV ping over a multi-segment pseudowire.

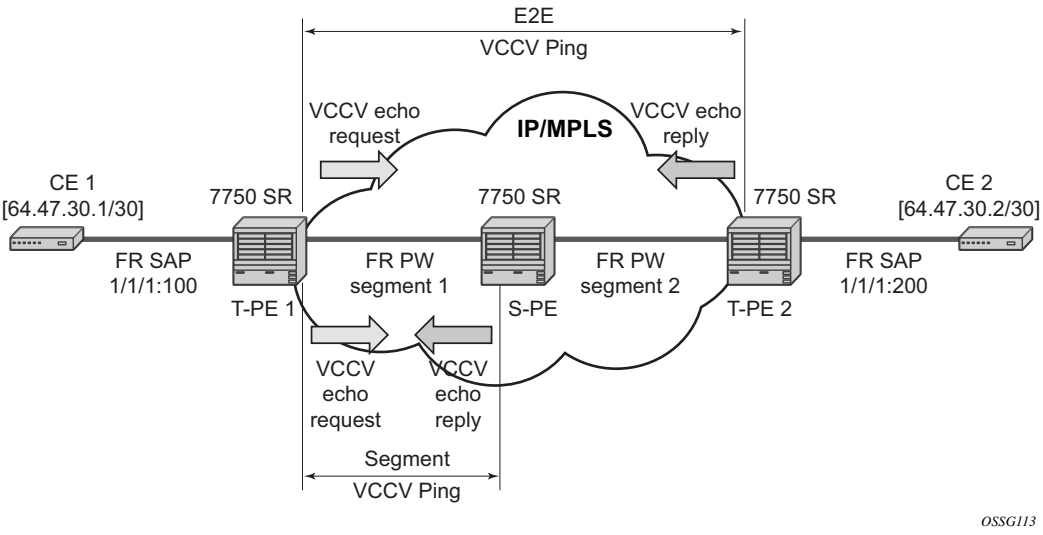

Figure 2: VCCV-Ping over a Multi-Segment Pseudowire

Pseudowire switching is a method to scale a large network of VLL or VPLS services by removing the need for a full mesh of T-LDP sessions between the 7x50 PE nodes as the number of these nodes grows over time. In the network displayed in Figure 72, a terminator PE (T-PE) is where the pseudowire originates and terminates. The switching PE (S-PE) is the node that performs switching by cross-connecting two spoke SDPs.

VCCV ping supports the following OAM functions:

- VCCV ping to a destination PE. A VLL FEC ping is a message sent by T-PE1 to test the FEC at T-PE2. The operation at T-PE1 and T-PE2 is the same as in the case of a single-segment pseudowire in Figure 72.

  An alternative method, based on T-PE1 including a new multi-segment pseudowire control word which gets processed by S-PE1, is described in draft-hart-pwe3-segmented-pw-vccv-00.txt, *VCCV Extensions for Segmented Pseudo-Wire*.

- Use of the OAM control word (Figure 73).

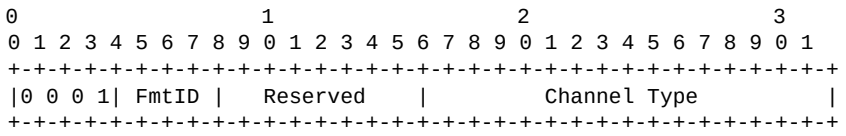

Figure 3: OAM Control Word Format

IGMP Snooping Diagnostics

MFIB Ping

The multicast forwarding information base (MFIB) ping OAM tool allows to easily verify inside a VPLS which SAPs would normally egress a certain multicast stream. The multicast stream is identified by a source unicast and destination multicast IP address, which are mandatory when issuing an MFIB ping command.

An MFIB ping packet will be sent through the data plane and goes out with the data plane format containing a configurable VC label TTL. This packet traverses each hop using forwarding plane information for next hop, VC label, etc. The VC label is swapped at each service-aware hop, and the VC TTL is decremented. If the VC TTL is decremented to 0, the packet is passed up to the management plane for processing. If the packet reaches an egress node, and would be forwarded out a customer facing port (SAP), it is identified by the OAM label below the VC label and passed to the management plane.

End-to-End Testing of Paths in an LDP ECMP Network

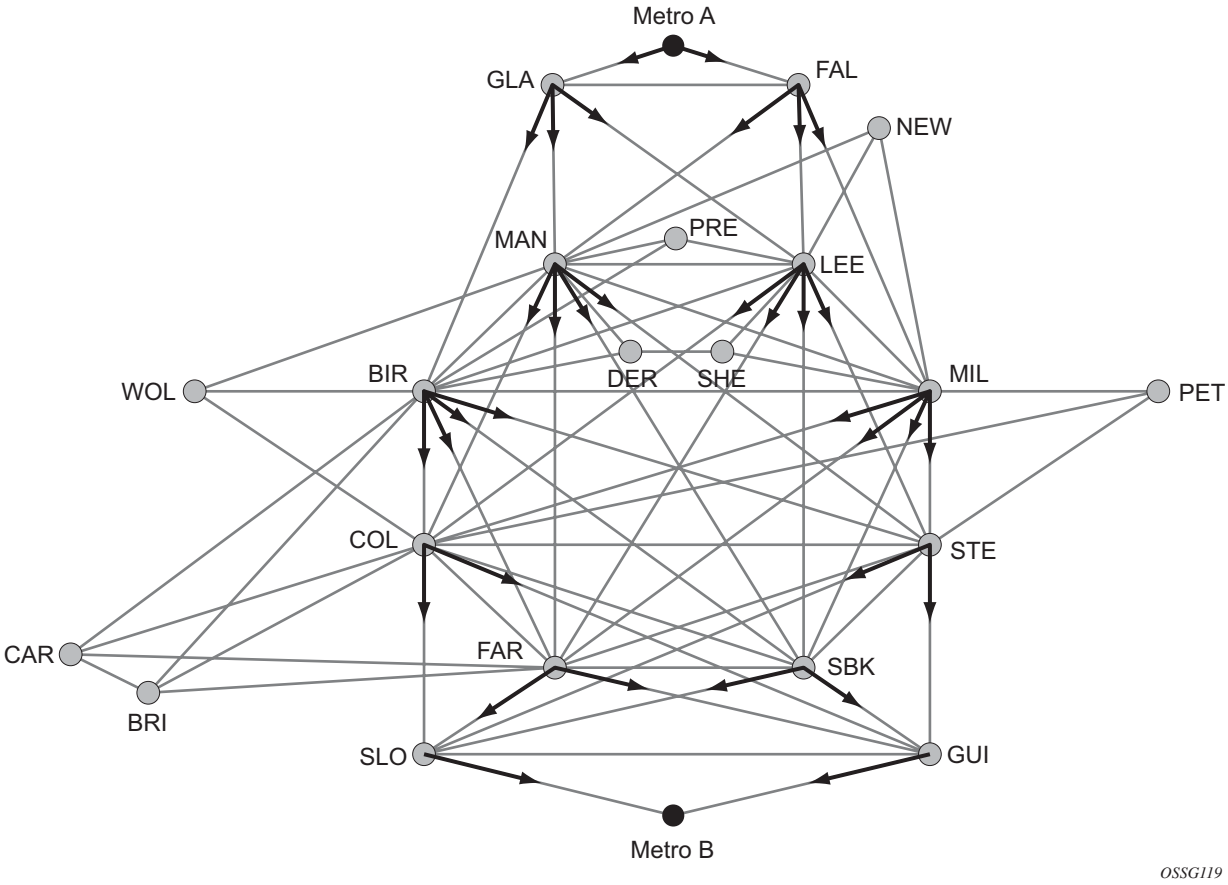

Figure 4: Network Resilience Using LDP ECMP

Figure 74 depicts faults that are detected through IGP and/or LDP are corrected as soon as IGP and LDP re-converge. The impacted traffic will be forwarded on the next available ECMP path as determined by the hash routine at the node that had a link failure.

However, there are faults which the IGP/LDP control planes may not detect. These faults are mainly due to a corruption of the control plane state or of the data plane state in a node. The LDP ECMP OAM is intended to detect these “silent” data plane and control plane faults. For example, it is possible that the forwarding plane of a node has a corrupt Next Hop Label Forwarding Entry (NHLFE) and keeps forwarding packets over an ECMP path only to have the downstream node discard them. This data plane fault can only be detected by an OAM tool that can test all possible end-to-end paths between the ingress LER and the egress LER. A corruption of the NLHFE entry can also result from a corruption in the control plane at that node.

Consider the following network topology:

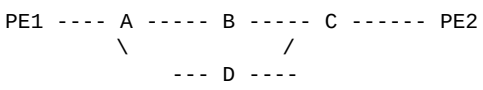

Assume LDP label distribution operates in the independent control mode and there is no active LDP session between nodes B and C. Node B will have an unlabeled entry for FEC PE2 and will forward the packet using the IGP path. The LDP ECMP OAM tool can detect the discontinuity of the LSP from PE1 to PE2.

Persistent loops are the other types of faults the LDP ECMP OAM can detect. Assume that link B-C failed and was subsequently restored. Node A can now forward to node B packets destined to PE2. A failure of the IGP/LDP re-convergence on link B-C will cause node B to forward those packets back to A since node B's NLHFE still indicates that the next hop for PE2 FEC is node A.

The LDP ECMP OAM feature consists of the two main capabilities:

1. An LDP ECMP tree building capability each ingress PE node will run to discover all possible paths to each egress PE FEC. This capability is based on having the ingress PE probe the network with LSP trace messages to collect the multipath information for a FEC from all nodes, including the egress LER.

   a. This tool can run periodically, anywhere between 15 minutes to an hour, on an ingress PE. There are parameters defined to stop the tool searching after a number of paths have been discovered for a FEC. Hence, the periodic mode of operation of the tool is not intended to cover 100% of the tree.

   b. Users can run the tool in a single run mode without bounds in order to discover the entire tree.

   c. When running in the periodic mode, the following are the user-configurable parameters to halt the probing process for a given egress FEC:
      → Maximum number of ECMP paths discovered per FEC.
      → Maximum number of hops into the network to probe for ECMP paths.

   d. There is no requirement to run this tool after a network event such as IGP re-convergence.

   e. The ingress PE automatically add an egress PE which FEC it learned to the list of probed destinations.

   f. Users can configure a policy to prevent specific FECs from being probed by the ingress PE.

g. Users can configure the values to write into the DSCP and the EXP fields of the LSP trace messages such that they get queued in the proper forwarding class queue in the network.

h. The processing of the LSP trace packets does not cause critical CPM/IOM tasks (for example, routing re-convergence) to be delayed.

2. A periodic path exercising capability to check the continuity of the discovered paths.

a. The continuity check tool uses LSP ping probes.

b. This tool can run periodically at a frequency of at least 1 probe per minute per discovered FEC path.

c. The processing of the probe messages does not cause critical CPM/IOM tasks (for example, routing re-convergence) to be delayed.

d. The ingress PE to can generate SNMP based alarms to the OSS layer to indicate faults conditions and clearing of faults. Note that failures which the network has recovered from does not cause alarms.

e. The alarms generated contain sufficient information to allow an OSS to localize the faulty node/link and to identify the service/customers impacted by the fault.

LDP ECMP Tree Building

The 7750 SR ingress LER builds the ECM tree for a given FEC (egress LER) by sending LSP trace messages and including the LDP IPv4 Prefix FEC TLV as well as the downstream mapping TLV.In order to build the ECMP tree, the 7750 SR LER inserts an IP address range drawn from the 127/8 space. When received by the downstream LSR, it will use this range to determine which ECMP path is exercised by any IP address or a sub-range of addresses within that range based on its internal hash routine. When the MPLS echo reply is received by the 7750 SR LER, it will record this information and proceed with the next echo request message targeted for a node downstream of the first LSR node along one of the ECMP paths. The sub-range of IP addresses indicated in the initial reply will be used since the objective is to have the LSR downstream of the 7750 SR LER pass this message to its downstream node along the first ECMP path.

The following figure illustrates the behavior through the following example adapted from RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*:

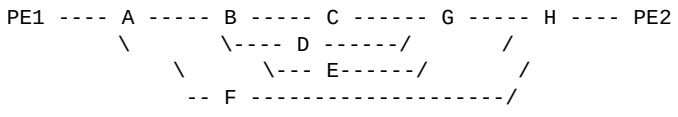

LSR A has two downstream LSRs, B and F, for PE2 FEC. PE1 receives an echo reply from A with the Multipath Type set to 4, with low/high IP addresses of 127.1.1.1->127.1.1.255 for downstream LSR B and 127.2.1.1->127.2.1.255 for downstream LSR F. PE1 reflects this information to LSR B. B, which has three downstream LSRs, C, D, and E, computes that 127.1.1.1->127.1.1.127 would go to C and 127.1.1.128-> 127.1.1.255 would go to D. B would then respond with 3 Downstream Mappings: to C, with Multipath Type 4 (127.1.1.1->127.1.1.127); to D, with Multipath Type 4 (127.1.1.127->127.1.1.255); and to E, with Multipath Type 0.

The 7750 SR supports multipath type 0 and 8, and up to a maximum of 36 bytes for the multipath length and supports the LER part of the LDP ECMP tree building feature.

A user configurable parameter sets the frequency of running the tree trace capability. The minimum and default value is 60 minutes and the increment is 1 hour.

The 7750 SR LER gets the list of FECs from the LDP FEC database. New FECs will be added to the discovery list at the next tree trace and not when they are learned and added into the FEC database. The maximum number of FECs to be discovered with the tree building feature is limited to 500. The user can configure FECs to exclude the use of a policy profile.

Periodic Path Exercising

The periodic path exercising runs in the background to test the LDP ECMP paths discovered by the tree building capability. The probe used is an LSP ping message with an IP address drawn from the sub-range of 127/8 addresses indicated by the output of the tree trace for this FEC.

The periodic LSP ping messages continuously probes an ECMP path at a user configurable rate of at least 1 message per minute. This is the minimum and default value. The increment is 1 minute. If an interface is down on a 7750 SR LER, then LSP ping probes that normally go out this interface will not be sent.

The LSP ping routine updates the content of the MPLS echo request message, specifically the IP address, as soon as the LDP ECMP tree trace has output the results of a new computation for the path in question.

Service Assurance Agent Overview

In the last few years, service delivery to customers has drastically changed. Services such as VPLS and VPRN are offered. The introduction of Broadband Service Termination Architecture (BSTA) applications such as Voice over IP (VoIP), TV delivery, video and high speed Internet services force carriers to produce services where the health and quality of Service Level Agreement (SLA) commitments are verifiable to the customer and internally within the carrier.

SAA is a feature that monitors network operations using statistics such as jitter, latency, response time, and packet loss. The information can be used to troubleshoot network problems, problem prevention, and network topology planning.

The results are saved in SNMP tables are queried by either the CLI or a management system. Threshold monitors allow for both rising and falling threshold events to alert the provider if SLA performance statistics deviate from the required parameters.

SAA Application

SAA allows two-way timing for several applications. This provides the carrier and their customers with data to verify that the SLA agreements are being properly enforced.

Two-way time measures requests from this node to the specified DNS server. This is done by performing an address request followed by an immediate release of the acquired address once the time measurement has been performed.

Traceroute Implementation

Various applications, such as lsp-trace, traceroute and vprn-trace, pass through the P-chip on the way to the control CPU. At this point, and when it egresses the control CPU, the P-chip should insert a timestamp inside the packet. Only packets processed by the Control CPU are processed.

When interpreting these timestamps care must be taken that some nodes are not capable of providing timestamps, as such timestamps must be associated with the same IP-address that is being returned to the originator to indicate what hop is being measured.

OAM/SAA List of Commands

Table 44 summarizes the commands and command uses. The command list is organized in the following task-oriented manner:

- LSP diagnostic commands
- SDP diagnostic commands
- Service diagnostic commands
- VPLS MAC diagnostic commands
- IGMP snooping diagnostic commands
- SAA configuration commands
- VLL diagnostic commands

Table 1: OAM Command Summary

Command	Description	Page
LSP diagnostic commands		
`lsp-ping`	In-band LSP ping utility to verify LSP connectivity.	1434
`lsp-trace`	In-band LSP traceroute command to determine the hop-by-hop path for an LSP.	1436
SDP diagnostic commands		
`sdp-mtu`	Performs in-band MTU Path tests on an SDP to determine the largest path-mtu supported on an SDP.	1401
`sdp-ping`	Tests an SDP for in-band uni-directional or round trip connectivity with a round trip time estimate.	1442
Service diagnostic commands		
`svc-ping`	Tests a service ID for correct and consistent provisioning between two service end points. The following information can be determined from svc-ping: • Local and remote service existence • Local and remote service state • Local and remote service type correlation • Local and remote customer association • Local and remote service-to-SDP bindings and state • Local and remote ingress and egress service label association	1403

Table 1: OAM Command Summary (Continued)

Command	Description	Page
VPLS MAC diagnostic commands		
mac-ping	In-band and out-of-band utility to determine the existence of an egress SAP binding of a given MAC within a VPLS. Utility can also be used to display all operationally up SAPs in the VPLS service.	1438
cpe-ping	In-band and out-of-band utility to determine connectivity to a given IP address within a VPLS.	
mac-populate	Populates the FIB with an OAM-type MAC entry indicating the node is the egress node for the MAC address and optionally floods the OAM MAC association throughout the service.	1418
mac-purge	Removes an OAM-type MAC entry from the FIB and optionally floods the OAM MAC removal throughout the service.	1421
mac-trace	In-band or out-of-band utility to determine the hop-by-hop path for a destination MAC address within a VPLS.	1440
IGMP snooping diagnostic commands		
mfib-ping	In-band utility to determine the list of SAPs which egress a certain IP multicast stream (identified by source unicast and destination multicast IP addresses) within a VPLS. Utility can also be used to display all SAPs which are operationally up in the VPLS service.	1422
SAA configuration commands		
config>saa>test		
description	Description for this SAA test.	1426
jitter-event	At the termination of an SAA test, both rising and falling thresholds are evaluated versus the configuration and events generated as required.	1427
latency-event	At the termination of an SAA test, both rising and falling thresholds are evaluated versus the configuration and events generated as required.	1427
loss-event	At the termination of an SAA test, both rising and falling thresholds are evaluated versus the configuration and events generated as required.	1428
shutdown/ no shutdown	In order to modify an existing test it must be shut down first. When a test is created it is in shutdown mode until a **no shutdown** command is executed.	1395
VLL diagnostic commands		
vccv-ping	Configures a Virtual Circuit Connectivity Verification (VCCV) test.	1449

Table 1: OAM Command Summary (Continued)

Command	Description	Page
SAA test type configuration commands		
config>saa>test>type		
cpe-ping	Configures a CPE ping test.	1429
dns	Configures a DNS name resolution test.	1431
icmp-ping	Specifies that icmp-ping packets be used for this test.	1431
icmp-trace	Configures an ICMP traceroute test.	1433
lsp-ping	Specifies that lsp-ping packets be used for this test.	1434
lsp-trace	Specifies that lsp-trace packets be used for this test.	1436
mac-ping	Specifies that mac-ping packets be used for this test.	1438
mac-trace	Specifies that mac-trace packets be used for this test.	1440
sdp-ping	Performs an SAA test on a SDP for either one-way or two-way timing.	1442
vccv-ping	Configures a VCCV ping test.	1449
vprn-ping	Performs an SAA VPRN-ping test to a specific node for one-way or optionally two-way timing.	1414
vprn-trace	Performs an SAA VPRN-trace test to a specific node for one-way or optionally two-way timing.	1415

Configuring SAA Test Parameters

Use the following CLI syntax to create SAA test parameters:

Example:
```
config# saa
        config>saa# test t1
        config>saa>test$ type
        config>saa>test>type$ lsp-ping to-104 interval 4 send-count 5
        config>saa>test>type$ exit
        config>saa>test# no shutdown
        config>saa>test# exit
        config>saa# exit
```

The following example displays the configuration:

```
A:ALA-48>config>saa# info
----------------------------------------------
        test "t1"
            type
                lsp-ping "to-104" interval 4 send-count 5
            exit
            no shutdown
        exit
----------------------------------------------
A:ALA-48>config>saa#
```

After running the test twice, the result is displayed below:

```
A:ALA-48>config>saa# show saa t1
Test Run: 1
Total number of attempts: 5
Number of requests that failed to be sent out: 1
Number of responses that were received: 4
Number of requests that did not receive any response: 0
Total number of failures: 1, Percentage: 20
Roundtrip Min: 0 ms, Max: 30 ms, Average: 15 ms, Jitter: 1 ms
Per test packet:
    Sequence: 1, Result: The active lsp-id is not found., Roundtrip: 0 ms
    Sequence: 2, Result: Response Received, Roundtrip: 0 ms
    Sequence: 3, Result: Response Received, Roundtrip: 0 ms
    Sequence: 4, Result: Response Received, Roundtrip: 30 ms
    Sequence: 5, Result: Response Received, Roundtrip: 30 ms
Test Run: 2
Total number of attempts: 5
Number of requests that failed to be sent out: 0
Number of responses that were received: 5
Number of requests that did not receive any response: 0
Total number of failures: 0, Percentage: 0
Roundtrip Min: 0 ms, Max: 40 ms, Average: 14 ms, Jitter: 5 ms
Per test packet:
    Sequence: 1, Result: Response Received, Roundtrip: 40 ms
    Sequence: 2, Result: Response Received, Roundtrip: 0 ms
    Sequence: 3, Result: Response Received, Roundtrip: 0 ms
    Sequence: 4, Result: Response Received, Roundtrip: 0 ms
    Sequence: 5, Result: Response Received, Roundtrip: 30 ms
```

OAM Command Reference

Command Hierarchies

Operational Commands

GLOBAL

- **ping** [*ip-address* | *dns-name*] [**rapid** | **detail**] [**ttl** *time-to-live*] [**tos** *type-of-service*] [**size** *bytes*] [**pattern** *pattern*] [**source** *ip-address* | *ipv6-address* | *dns-name*] [**interval** *seconds*] [{**next-hop** *ip-address*} | {**interface** *interface-name*} | **bypass-routing**] [**count** *requests*] [**do-not-fragment**] [**router** *router-instance*] [**timeout** *timeout*]
- **traceroute** [*ip-address* | *dns-name*] [**ttl** *ttl*] [**wait** *milli-seconds*] [**no-dns**][**source** *src-ip-address*] [**tos** *type-of-service*] [**router** [*router-instance*]

ATM Diagnostics

GLOBAL

- **oam**
  - **atm-ping** *port-id*:*vpi*/*vci* [**end-to-end** | **segment**] [**dest** *destination-id*][**send-count** *send-count*][**timeout** *seconds*][**interval** *seconds*]

LDP Diagnostics

GLOBAL

- **oam**
  - **ldp-treetrace** {**prefix** *ip-prefix/mask*} [**max-ttl** *ttl-value*] [**max-path** *max-paths*] [**timeout** *timeout*] [**retry-count** *retry-count*] [**fc** *fc-name* [**profile** *profile*]]
- **config**
  - **test-oam**
    - [**no**] **ldp-treetrace**
      - **fc** *fc-name* [profile {in|out}]
      - **no fc**
      - **path-discovery**
        - **interval** *minutes*
        - **no interval**
        - **max-path** *max-paths*
        - **no max-path**
        - **max-ttl** *ttl-value*
        - **no max-ttl**
        - **policy-statement** *policy-name*[...(up to 5 max)]
        - **no policy-statement**
        - **retry-count** *retry-count*
        - **no retry-count**
        - **timeout** *timeout*
        - **no timeout**
      - **path-probing**
        - **interval** *minutes*
        - **no interval**
        - **retry-count** *retry-count*
        - **no retry-count**
        - **timeout** *timeout*
        - **no timeout**
      - [**no**] **shutdown**

LSP Diagnostics

GLOBAL

— **oam**

— **lsp-ping** {{[*lsp-name]* [**path** path-name]} | {**prefix** ip-prefix/mask}} [**fc** *fc-name*] [**profile** {**in** | **out**}]] [**size** *octets*] [**ttl** *label-ttl*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**path-destination** *ip-address*] [**interface** *if-name* | **next-hop** *ip-address*]][**detail**]

— **lsp-trace** {{[*lsp-name]* [**path** path-name]} | {**prefix** ip-prefix/mask}} [**fc** *fc-name*] [**profile** {**in** | **out**}]] [**max-fail** *no-response-count*] [**probe-count** *probes-per-hop*] [**size** *octets*][**min-ttl** *min-label-ttl*]] [**max-ttl** *max-label-ttl*] [**timeout** *timeout*] [[**interval** *interval*] [**path-destination** *ip-address*] [**interface** *if-name* | **next-hop** *ip-address*]][**detail**]

SDP Diagnostics

GLOBAL

— **oam**

— **sdp-mtu** *orig-sdp-id* **size-inc** *start-octets end-octets* [**step** *step-size*] [**timeout** *seconds*] [**interval** *seconds*]

— **sdp-ping** *orig-sdp-id* [**resp-sdp** *resp-sdp-id*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**timeout** *seconds*] [**interval** *seconds*] [**size** *octets*] [**count** *send-count*]

Service Diagnostics

GLOBAL

— **oam**

— **ancp** {**subscriber** *sub-ident-string* | **ancp-string** *ancp-string*} **loopback** [**count** *count*] [**timeout** *seconds*] [**alarm**]

— **ancp** **subscriber** *sub-ident-string* **loopback** [**count** *send-count*] [**timeout** *seconds*] [**alarm**]

— **svc-ping** {*ip-addr* | *dns-name*} **service** *service-id* [**local-sdp**] [**remote-sdp**]

— **host-connectivity-verify** **service** *service-id* [**sap** *sap-id*]

— **host-connectivity-verify** **subscriber** *sub-ident-string* [**sla-profile** *sla-profile-name*]

— **vprn-ping** *service-id* **source** *src-ip* **destination** *ip-address* [**fc** *fc-name* [**profile** {**in** | **out**}] [**size** *size*] [**ttl** *vc-label-ttl*] [**return-control**] [**interval** *interval*] [**count** *send-count*] [**timeout** *timeout*]

— **vprn-trace** *service-id* **source** *src-ip* **destination** *ip-address* [**fc** *fc-name* [**profile** {**in** | **out**}] [**size** *size*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**return-control**] [**probe-count** *send-count*] [**interval** *seconds*] [**timeout** *timeout*]

VLL Diagnostics

GLOBAL

— **oam**

— **vccv-ping** *sdp-id:vc-id* [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr* **pw-id** *pw-id*][**reply-mode** {**ip-routed** | **control-channel**}][**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*][**ttl** *vc-label-ttl*]

VPLS MAC Diagnostics

GLOBAL

— **oam**

— **cpe-ping** **service** *service-id* **destination** *dst-ieee-address* **source** *ip-address* [**source-mac** *ieee-address*][**ttl** *vc-label-ttl*] [**count** *send-count*] [**send-control**] [**return-control**] [**interval** *interval*]

— **dns** **target-addr** *dns-name* **name-server** *ip-address* [**source** *ip-address*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]

— **mac-ping** **service** *service-id* **destination** *dst-ieee-address* [**source** *src-ieee-address*] [**fc** *fc-name* [**profile** **in** | **out**]] [**size** *octets*] [**ttl** *vc-label-ttl*] [**count** *send-count*] [**send-control**] [**return-control**] [**interval** *interval*] [**timeout** *timeout*]

- **mac-populate** *service-id* **mac** *ieee-address* [**flood**] [**age** *seconds*] [**force**] [**target-sap** *sap-id*] [**send-control**]
- **mac-purge** *service-id* **target** *ieee-address* [**flood**] [**send-control**] [**register**]
- **mac-trace** *service-id* **destination** *ieee-address* [**size** *octets*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**send-control**] [**return-control**] [**source** *ieee-address*] [**probe-count** *probes-per-hop*] [**interval** *interval*]
- **mac-trace service** *service-id* **destination** *ieee-address* [**source** *ieee-address*] [**fc** *fc-name* [**profile in**|**out**]] [**size** *octets*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**probe-count** *send-count*] [**send-control**] [**return-control**] [**interval** *interval*] [**timeout** *timeout*]
- **mfib-ping service** *service-id* **source** *src-ip* **destination** *mcast-address* [**size** *size*] [**ttl** *vc-label-ttl*] [**return-control**] [**interval** *interval*] [**count** *send-count*] [**timeout** *timeout*]

Ethernet in the First Mile (EFM) Commands

- **efm** *port-id*
  - **local-loopback** {**start** | **stop**}
  - **remote-loopback** {**start** | **stop**}

SAA Command Reference

Command Hierarchies

Configuration Commands

```
config
    — saa
        — [no] test test-name [owner test-owner]
            — description description-string
            — no description
            — [no] jitter-event rising-threshold threshold [falling-threshold threshold] [direction]
            — [no] latency-event rising-threshold threshold [falling-threshold threshold] [direction]
            — [no] loss-event  rising-threshold threshold [falling-threshold threshold]  [direction]
            — [no] shutdown
            — [no] type
                — cpe-ping service service-id destination ip-address source ip-address [source-mac ieee-address] [fc fc-name [profile [in | out]][ttl vc-label-ttl] [count send-count] [send-control] [return-control] [interval interval]
                — dns target-addr dns-name name-server ip-address [source  ip-address] [count send-count] [timeout timeout] [interval interval]
                — icmp-ping [ip-address | dns-name] [rapid | detail] [ttl time-to-live] [tos type-of-service] [size bytes] [pattern pattern] [source ip-address | dns-name] [interval seconds] [{next-hop ip-address}|{interface interface-name}|bypass-routing] [count requests] [do-not-fragment] [router router-instance] [timeout timeout]
                — icmp-trace [ip-address | dns-name] [ttl time-to-live] [wait milli-seconds] [tos type-of-service] [source ip-address] [tos type-of-service] [router router-instance]
                — lsp-ping {{lsp-name [path path-name]}|{prefix ip-prefix/mask}} [size octets] [ttl label-ttl] [timeout timeout] [interval interval] [fc {be|l2|af|l1|h2|ef|h1|nc} [profile {in|out}]] [send-count send-count]
                — lsp-trace {lsp-name [path path-name]} | {prefix ip-prefix/mask}} [size octets] [min-ttl min-label-ttl] [max-ttl max-label-ttl] [max-fail no-response-count] [send-count send-count] [timeout timeout] [interval interval] [fc fc-name [profile {in | out}]]
                — mac-ping service service-id destination ieee-address [size octets] [ttl vc-label-ttl] [send-control] [return-control] [source ieee-address] [interval interval] [count send-count]
                — mac-trace  service service-id destination ieee-address [size octets] [min-ttl vc-label-ttl] [max-ttl vc-label-ttl] [send-control] [return-control] [source ieee-address] [probe-count probes-per-hop] [interval interval]
                — sdp-ping orig-sdp-id [resp-sdp resp-sdp-id] [fc fc-name [profile {in | out}]] [timeout seconds] [interval seconds] [size octets] [count send-count]
                — vccv-ping sdp-id:vc-id [src-ip-address ip-addr dst-ip-address ip-addr pw-id pw-id][reply-mode {ip-routed | control-channel}][fc fc-name
```

[**profile** {**in** | **out**}]] [**size** *octets*] [**count** *send-count*][**timeout** *timeout*] [**interval** *interval*][**ttl** *vc-label-ttl*]

- — **vprn-ping** *service-id* source *src-ip* **destination** *dst-ip* [**size** *size*] [**ttl** *vc-label-ttl*] [**return-control**] [**interval** *interval*] [**count** *send-count*] [**timeout** *timeout*]
- — **vprn-trace** *service-id* source *src-ip* **destination** *dst-ip* [**size** *size*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**return-control**] [**probe-count** *probes-per-hop*] [**interval** *seconds*] [**timeout** *timeout*]

SAA Diagnostics

GLOBAL

— oam

— **saa** *test-name* [**owner** *test-owner*] {**start**|**stop**}

Show Commands

show
- **saa** [*test-name* [**owner** *test-owner*]]
- **test-oam**
  - **ldp-treetrace** [**prefix** *ip-prefix/mask*] [**detail**]

Clear Commands

clear
- **saa** [*test-name* [**owner** *test-owner*]]

Debug Commands

debug
- [**no**] **lsp-ping-trace**
- [**no**] **oam**
  - **lsp-ping-trace** [**tx** | **rx** | **both**] [**raw** | **detail**]
  - **no lsp-ping-trace**

OAM and SAA Commands

Command Hierarchies

Operational Commands

shutdown

Syntax [**no**] **shutdown**

Context config>saa>test

In order to modify an existing test it must first be shut down. When a test is created it will be in shutdown mode until a **no shutdown** command is executed.

A **shutdown** can only be performed if a test is not executing at the time the command is entered.

Use the **no** form of the command to set the state of the test to operational.

shutdown

Syntax [**no**] **shutdown**

Context config>test-oam>ldp-treetrace

Description This command suspends the background process running the LDP ECMP OAM tree discovery and path probing features. The configuration is not deleted.

Use the **no** form of the command to enable the background process.

ping

Syntax **ping** [*ip-address* | *dns-name*] [**rapid** | **detail**] [**ttl** *time-to-live*] [**tos** *type-of-service*] [**size** *bytes*] [**pattern** *pattern*] [**source** *ip-address* | *dns-name*] [**interval** *seconds*] [{**next-hop** *ip-address*} | {**interface** *interface-name*} | **bypass-routing**] [**count** *requests*] [**do-not-fragment**] [**router** *router-instance*] [**timeout** *timeout*]

Context <GLOBAL>

Description This command verifies the reachability of a remote host.

Parameters *ip-address* — The far-end IP address to which to send the **svc-ping** request message in dotted decimal notation.

Values

ipv4-address:	a.b.c.d	
ipv6-address:	x:x:x:x:x:x:x:x	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 — FFFF]H
	d:	[0 — 255]D

dns-name — The DNS name of the far-end device to which to send the **svc-ping** request message, expressed as a character string.

Values 1 - 2147483647

rapid — Packets will be generated as fast as possible instead of the default 1 per second.

detail — Displays detailed information.

ttl *time-to-live* **—** The TTL value for the MPLS label, expressed as a decimal integer.

Values 1 — 128

tos *type-of-service* **—** Specifies the service type.

Values 0 — 255

size *bytes* **—** The request packet size in bytes, expressed as a decimal integer.

Values 0 — 16384

pattern ***pattern*** **—** The date portion in a ping packet will be filled with the pattern value specified. If not specified, position info will be filled instead.

Values 0 — 65535

source *ip-address* **—** Specifies the IP address to be used.

Values

ipv4-address:	a.b.c.d	
ipv6-address:	x:x:x:x:x:x:x:x	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 — FFFF]H
	d:	[0 — 255]D

router router-instance **—** Specifies the router name or service ID.

Values

router-name:	Base , management
service-id:	1 — 2147483647

Default Base

bypass-routing — Specifies whether to send the ping request to a host on a directly attached network bypassing the routing table.

interface *interface-name* — Specifies the name of an IP interface. The name must already exist in the **config>router>interface** context.

next-hop *ip-address* **—** Only displays static routes with the specified next hop IP address.

Values

ipv4-address:	a.b.c.d (host bits must be 0)
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d

x: [0 — FFFF]H
d: [0 — 255]D

count *requests* **—** Specifies the number of times to perform an OAM ping probe operation. Each OAM echo message request must either timeout or receive a reply before the next message request is sent.

Values 1 — 100000

Default 5

do-not-fragment — Sets the DF (Do Not Fragment) bit in the ICMP ping packet.

timeout *seconds* **—** Overrides the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A 'request timeout' message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

traceroute

Syntax **traceroute** [*ip-address* | *dns-name*] [**ttl** *ttl*] [**wait** *milli-seconds*] [**no-dns**] [**source** *ip-address*] [**tos** *type-of-service*] [**router** *router-instance*]

Context <GLOBAL>

Description The TCP/IP traceroute utility determines the route to a destination address. DNS lookups of the responding hosts is enabled by default.

```
*A:ALA-1# traceroute 192.168.xx.xx4
traceroute to 192.168.xx.xx4, 30 hops max, 40 byte packets
 1  192.168.xx.xx4 0.000 ms  0.000 ms  0.000 ms
*A:ALA-1#
```

Parameters *ip-address* — The far-end IP address to which to send the traceroute request message in dotted decimal notation.

Values
ipv4-address : a.b.c.d
ipv6-address: x:x:x:x:x:x:x:x
x:x:x:x:x:x:d.d.d.d
x: [0 — FFFF]H
d: [0 — 255]D

dns-name — The DNS name of the far-end device to which to send the traceroute request message, expressed as a character string.

ttl *ttl* — The maximum Time-To-Live (TTL) value to include in the traceroute request, expressed as a decimal integer.

Values 1 — 255

wait *milliseconds* **—** The time in milliseconds to wait for a response to a probe, expressed as a decimal integer.

Default 5000

Values 1 — 60000

no-dns — When the **no-dns** keyword is specified, DNS lookups of the responding hosts will not be performed, only the IP addresses will be printed.

Default DNS lookups are performed

source *ip-address* **—** The source IP address to use as the source of the probe packets in dotted decimal notation. If the IP address is not one of the device's interfaces, an error is returned.

tos *type-of-service* **—** The type-of-service (TOS) bits in the IP header of the probe packets, expressed as a decimal integer.

Values 0 — 255

router *router-name* **—** Specify the alphanumeric character string up to 32 characters.

Default Base

router *service-id* **—** The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR on which this service is defined.

Values 1 — 2147483647

ATM Diagnostics

atm-ping

Syntax **atm-ping** *port-id*: *vpi/vci* [**end-to-end** | **segment**] [**dest** *destination-id*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *seconds*]

Context <GLOBAL>

Description This command tests ATM path connectivity and round trip time on an ATM VCC.

Parameters *port-id:vpi/vci* — Specifies the ID of the access port of the target VC. This parameter is required.

Values

port-id	*slot/mda/port*
aps-id	aps-*group-id*
	aps keyword
	group-id 1 — 64
vpi	0 — 4095 (NNI)
	0 — 255 (UNI)
vci	1, 2, 5 — 65535

end-to-end | **segment —** Specifies whether the ATM OAM loopback cell is destined to the first segment point in the line direction or the PVCC's connection endpoint.

dest *destination-id* **—** Defines the LLID field in an OAM loopback cell. If set to all 1s, only the connection end (end-to-end ping) or segment end (segment ping) will respond to the ping. If the 'segment' parameter is specified and 'dest' is set to a specific destination, only the destination will respond to the ping.

Values A 16 byte octet string, with each octet separated by a colon, if not specified the value of all 0x11 will be used.

send-count *send-count* **—** The number of messages to send, expressed as a decimal integer. The **send-count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

Service Diagnostics

ancp

Syntax **ancp** {**subscriber** *sub-ident-string* | **ancp-string** *ancp-string*} **loopback** [**count** *count*] [**timeout** *seconds*] [**alarm**]
ancp subscriber *sub-ident-string* **loopback** [**count** *send-count*] [**timeout** *seconds*] [**alarm**]

Context <global>

Description This command sends an OAM request to the access node. ANCP can be used to send OAM messages to the access node. The access node must be able to accept these messages and will signal such support by the capability negotiations. If the operator attempts to send an OAM command to an access node that does not support such command the operation results in an error.

Parameters **subscriber** ***sub-ident-string*** **—** Specifies an existing subscriber-id. The node will use the ancp-string associated with the provided subscriber-id to identify the circuit.

ancp-string *ancp-string* **—** Specifies an existing ANCP string.

count *send-count* **—** Specifies the number of messages the access node will use to test the circuit. If omitted, the number will be determined by the access node via local policy.

1 — 32

timeout *seconds* **—** Specifies how long the controlling node will wait for a result.

0 — 300

alarm — Specifies that the CLI the result will be retuned to the CLI and a trap will be issued to indicate the test finished. If the flag is used through SNMP the results will be available in the results MIB and after the node sent the trap to indicate the results are ready.

loopback — Sendsan OAM loopback test request to the access node

sdp-mtu

Syntax **sdp-mtu** *orig-sdp-id* **size-inc** *start-octets end-octets* [**step** *step-size*] [**timeout** *seconds*] [**interval** *seconds*]

Context <GLOBAL>

Description Performs MTU Path tests on an SDP to determine the largest path-mtu supported on an SDP.

The **size-inc** parameter can be used to easily determine the **path-mtu** of a given SDP-ID. The forwarding class is assumed to be Best-Effort Out-of-Profile. The message reply is returned with IP/GRE encapsulation from the far-end 7750 SR. OAM request messages sent within an IP/GRE SDP must have the 'DF' IP header bit set to 1 to prevent message fragmentation.

To terminate an **sdp-mtu** in progress, use the CLI break sequence <Ctrl-C>.

Special Cases

SDP Path MTU Tests — SDP Path MTU tests can be performed using the **sdp-mtu size-inc** keyword to easily determine the **path-mtu** of a given SDP-ID. The forwarding class is assumed to be Best-Effort Out-of-Profile. The message reply is returned with IP/GRE encapsulation from the far-end 7750 SR.

With each OAM Echo Request sent using the **size-inc** parameter, a response line is displayed as message output. The path MTU test displays incrementing packet sizes, the number sent at each size until a reply is received and the response message.

As the request message is sent, its size value is displayed followed by a period for each request sent of that size. Up to three requests will be sent unless a valid response is received for one of the requests at that size. Once a response is received, the next size message is sent.

The response message indicates the result of the message request.

After the last reply has been received or response timeout, the maximum size message replied to indicates the largest size OAM Request message that received a valid reply.

Parameters

orig-sdp-id — The SDP-ID to be used by **sdp-ping,** expressed as a decimal integer. The far-end address of the specified SDP-ID is the expected *responder-id* within each reply received. The specified SDP-ID defines the encapsulation of the SDP tunnel encapsulation used to reach the far end. This can be IP/GRE or MPLS. If *orig-sdp-id* is invalid or administratively down or unavailable for some reason, the SDP Echo Request message is not sent and an appropriate error message is displayed (once the **interval** timer expires, sdp-ping will attempt to send the next request if required).

Values 1 — 17407

size-inc *start-octets end-octets* **—** Indicates an incremental Path MTU test will be performed with by sending a series of message requests with increasing MTU sizes. The *start-octets* and *end-octets* parameters are described below.

start-octets — The beginning size in octets of the first message sent for an incremental MTU test, expressed as a decimal integer.

Values 40 — 9198

end-octets — The ending size in octets of the last message sent for an incremental MTU test, expressed as a decimal integer. The specified value must be greater than *start-octets*.

Values 40 — 9198

step *step-size* **—** The number of octets to increment the message size request for each message sent for an incremental MTU test, expressed as a decimal integer. The next size message will not be sent until a reply is received or three messages have timed out at the current size.

If the incremented size exceeds the *end-octets* value, no more messages will be sent.

Default 32

Values 1 — 512

timeout *seconds* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A ‘request timeout’ message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

interval *seconds* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

Output **Sample SDP MTU Path Test Sample Output**

```
*A:router 1> sdp-mtu 6 size-inc 512 3072 step 256
   Size        Sent      Response
 -------       ----      --------------------------
   512         .         Success
   768         .         Success
  1024         .         Success
  1280         .         Success
  1536         .         Success
  1792         .         Success
  2048         .         Success
  2304         …         Request Timeout
  2560         …         Request Timeout
  2816         …         Request Timeout
  3072         …         Request Timeout
Maximum Response Size:   2048
```

svc-ping

Syntax **svc-ping** *ip-address* [**service** *service-id*] [**local-sdp**] [**remote-sdp**]

Context <GLOBAL>

Description Tests a service ID for correct and consistent provisioning between two service end points.

The **svc-ping** command accepts a far-end IP address and a Service-ID for local and remote service testing. The following information can be determined from **svc-ping**:

1. Local and remote service existence
2. Local and remote service state
3. Local and remote service type correlation
4. Local and remote customer association
5. Local and remote service-to-SDP bindings and state
6. Local and remote ingress and egress service label association

Unlike **sdp-ping**, only a single message will be sent per command; no count nor interval parameter is supported and round trip time is not calculated. A timeout value of 10 seconds is used before failing the request. The forwarding class is assumed to be Best-Effort Out-of-Profile

If no request is sent or a reply is not received, all remote information will be shown as N/A.

To terminate a **svc-ping** in progress, use the CLI break sequence <Ctrl-C>.

Upon request timeout, message response, request termination, or request error the following local and remote information will be displayed. Local and remote information will be dependent upon service existence and reception of reply.

Field	Description	Values
Request Result	The result of the **svc-ping** request message.	Sent - Request Timeout
		Sent - Request Terminated
		Sent - Reply Received
		Not Sent - Non-Existent Service-ID
		Not Sent - Non-Existent SDP for Service
		Not Sent - SDP For Service Down
		Not Sent - Non-existent Service Egress Label
Service-ID	The Service-ID being tested.	*service-id*
Local Service Type	The type of service being tested. If *service-id* does not exist locally, N/A is displayed.	Epipe, Ipipe, Fpipe, Apipe
		TLS
		IES
		Mirror-Dest
		N/A
Local Service Admin State	The local administrative state of *service-id*. If the service does not exist locally, the administrative state will be Non-Existent.	Admin-Up
		Admin-Down
		Non-Existent
Local Service Oper State	The local operational state of *service-id*. If the service does not exist locally, the state will be N/A.	Oper-Up
		Oper-Down
		N/A

Field	Description	Values (Continued)
Remote Service Type	The remote type of service being tested. If *service-id* does not exist remotely, N/A is displayed.	Epipe, Ipipe, Fpipe, Apipe
		TLS
		IES
		Mirror-Dest
		N/A
Remote Service Admin State	The remote administrative state of *service-id*. If the service does not exist remotely, the administrative state is Non-Existent.	Up
		Down
		Non-Existent
Local Service MTU	The local **service-mtu** for *service-id*. If the service does not exist, N/A is displayed.	*service-mtu*
		N/A
Remote Service MTU	The remote **service-mtu** for *service-id*. If the service does not exist remotely, N/A is displayed.	*remote-service-mtu*
		N/A
Local Customer ID	The local *customer-id* associated with *service-id*. If the service does not exist locally, N/A is displayed.	*customer-id*
		N/A
Remote Customer ID	The remote *customer-id* associated with *service-id*. If the service does not exist remotely, N/A is displayed.	*customer-id*
		N/A
Local Service IP Address	The local system IP address used to terminate remotely configured SDP-ID (as the **far-end** address). If an IP interface has not been configured to be the system IP address, N/A is displayed.	*system-ip-address*
		N/A
Local Service IP Interface Name	The name of the local system IP interface. If the local system IP interface has not been created, N/A is displayed.	*system-interface-name*
		N/A
Local Service IP Interface State	The state of the local system IP interface. If the local system IP interface has not been created, Non-Existent is displayed.	Up
		Down
		Non-Existent
Expected Far-end Address	The expected IP address for the remote system IP interface. This must be the **far-end** address entered for the **svc-ping** command.	*orig-sdp-far-end-addr*
		dest-ip-addr
		N/A
Actual Far-end Address	The returned remote IP address. If a response is not received, the displayed value is N/A. If the far-end service IP interface is down or non-existent, a message reply is not expected. **sdp-ping** should also fail.	*resp-ip-addr*
		N/A

Field	Description	Values (Continued)
Responders Expected Far-end Address	The expected source of the originator's SDP-ID from the perspective of the remote 7750 SR terminating the SDP-ID. If the far-end cannot detect the expected source of the ingress SDP-ID or the request is transmitted outside the SDP-ID, N/A is displayed.	*resp-rec-tunnel-far-end-address*
		N/A
Originating SDP-ID	The SDP-ID used to reach the **far-end** IP address if **sdp-path** is defined. The originating SDP-ID must be bound to the *service-id* and terminate on the **far-end** IP address. If an appropriate originating SDP-ID is not found, Non-Existent is displayed.	orig-sdp-id
		Non-Existent
Originating SDP-ID Path Used	Whether the Originating 7750 SR used the originating SDP-ID to send the **svc-ping** request. If a valid originating SDP-ID is found, operational and has a valid egress service label, the originating 7750 SR should use the SDP-ID as the requesting path if **sdp-path** has been defined. If the originating 7750 SR uses the originating SDP-ID as the request path, Yes is displayed. If the originating 7750 SR does not use the originating SDP-ID as the request path, No is displayed. If the originating SDP-ID is non-existent, N/A is displayed.	Yes
		No
		N/A
Originating SDP-ID Administrative State	The local administrative state of the originating SDP-ID. If the SDP-ID has been shutdown, Admin-Down is displayed. If the originating SDP-ID is in the no shutdown state, Admin-Up is displayed. If an originating SDP-ID is not found, N/A is displayed.	Admin-Up
		Admin-Up
		N/A
Originating SDP-ID Operating State	The local operational state of the originating SDP-ID. If an originating SDP-ID is not found, N/A is displayed.	Oper-Up
		Oper-Down
		N/A
Originating SDP-ID Binding Admin State	The local administrative state of the originating SDP-IDs binding to *service-id*. If an SDP-ID is not bound to the service, N/A is displayed.	Admin-Up
		Admin-Up
		N/A
Originating SDP-ID Binding Oper State	The local operational state of the originating SDP-IDs binding to *service-id*. If an SDP-ID is not bound to the service, N/A is displayed.	Oper-Up
		Oper-Down
		N/A
Responding SDP-ID	The SDP-ID used by the far end to respond to the **svc-ping** request. If the request was received without the **sdp-path** parameter, the responding 7750 SR will not use an SDP-ID as the return path, but the appropriate responding SDP-ID will be displayed. If a valid SDP-ID return path is not found to the originating 7750 SR that is bound to the *service-id*, Non-Existent is displayed.	*resp-sdp-id*
		Non-Existent

Field	Description	Values (Continued)
Responding SDP-ID Path Used	Whether the responding 7750 SR used the responding SDP-ID to respond to the **svc-ping** request. If the request was received via the originating SDP-ID and a valid return SDP-ID is found, operational and has a valid egress service label, the far-end 7750 SR should use the SDP-ID as the return SDP-ID. If the far end uses the responding SDP-ID as the return path, Yes is displayed. If the far end does not use the responding SDP-ID as the return path, No is displayed. If the responding SDP-ID is non-existent, N/A is displayed.	Yes No N/A
Responding SDP-ID Administrative State	The administrative state of the far-end SDP-ID associated with the return path for *service-id*. When a return path is administratively down, Admin-Down is displayed. If the return SDP-ID is administratively up, Admin-Up is displayed. If the responding SDP-ID is non-existent, N/A is displayed.	Admin-Up Admin-Up N/A
Responding SDP-ID Operational State	The operational state of the far-end SDP-ID associated with the return path for *service-id*. When a return path is operationally down, Oper-Down is displayed. If the return SDP-ID is operationally up, Oper-Up is displayed. If the responding SDP-ID is non-existent, N/A is displayed.	Oper-Up Oper-Down N/A
Responding SDP-ID Binding Admin State	The local administrative state of the responder's SDP-ID binding to *service-id*. If an SDP-ID is not bound to the service, N/A is displayed.	Admin-Up Admin-Down N/A
Responding SDP-ID Binding Oper State	The local operational state of the responder's SDP-ID binding to *service-id*. If an SDP-ID is not bound to the service, N/A is displayed.	Oper-Up Oper-Down N/A
Originating VC-ID	The Originators VC-ID associated with the SDP-ID to the far-end address that is bound to *service-id*. If the SDP-ID signaling is off, *originator-vc-id* is 0. If the *originator-vc-id* does not exist, N/A is displayed.	*originator-vc-id* N/A
Responding VC-ID	The Responder's VC-ID associated with the SDP-ID to *originator-id* that is bound to *service-id*. If the SDP-ID signaling is off or the service binding to SDP-ID does not exist, *responder-vc-id* is 0. If a response is not received, N/A is displayed.	*responder-vc-id* N/A
Originating Egress Service Label	The originating service label (VC-Label) associated with the *service-id* for the originating SDP-ID. If *service-id* does not exist locally, N/A is displayed. If *service-id* exists, but the egress service label has not been assigned, Non-Existent is displayed.	*egress-vc-label* N/A Non-Existent

Field	Description	Values (Continued)
Originating Egress Service Label Source	The originating egress service label source. If the displayed egress service label is manually defined, Manual is displayed. If the egress service label is dynamically signaled, Signaled is displayed. If the *service-id* does not exist or the egress service label is non-existent, N/A is displayed.	Manual Signaled N/A
Originating Egress Service Label State	The originating egress service label state. If the originating 7750 SR considers the displayed egress service label operational, Up is displayed. If the originating 7750 SR considers the egress service label inoperative, Down is displayed. If the *service-id* does not exist or the egress service label is non-existent, N/A is displayed.	Up Down N/A
Responding Service Label	The actual responding service label in use by the far-end 7750 SR for this *service-id* to the originating 7750 SR. If *service-id* does not exist in the remote 7750 SR, N/A is displayed. If *service-id* does exist remotely but the remote egress service label has not been assigned, Non-Existent is displayed.	*rec-vc-label* N/A Non-Existent
Responding Egress Service Label Source	The responder's egress service label source. If the responder's egress service label is manually defined, Manual is displayed. If the responder's egress service label is dynamically signaled, Signaled is displayed. If the *service-id* does not exist on the responder or the responder's egress service label is non-existent, N/A is displayed.	Manual Signaled N/A
Responding Service Label State	The responding egress service label state. If the responding 7750 SR considers its egress service label operational, Up is displayed. If the responding 7750 SR considers its egress service label inoperative, Down is displayed. If the *service-id* does not exist or the responder's egress service label is non-existent, N/A is displayed.	Up Down N/A
Expected Ingress Service Label	The locally assigned ingress service label. This is the service label that the far-end is expected to use for *service-id* when sending to the originating 7750 SR. If *service-id* does not exist locally, N/A is displayed. If *service-id* exists but an ingress service label has not been assigned, Non-Existent is displayed.	*ingress-vc-label* N/A Non-Existent
Expected Ingress Label Source	The originator's ingress service label source. If the originator's ingress service label is manually defined, Manual is displayed. If the originator's ingress service label is dynamically signaled, Signaled is displayed. If the *service-id* does not exist on the originator or the originators ingress service label has not been assigned, N/A is displayed.	Manual Signaled N/A
Expected Ingress Service Label State	The originator's ingress service label state. If the originating 7750 SR considers its ingress service label operational, Up is displayed. If the originating 7750 SR considers its ingress service label inoperative, Down is displayed. If the *service-id* does not exist locally, N/A is displayed.	Up Down N/A

Field	Description	Values (Continued)
Responders Ingress Service Label	The assigned ingress service label on the remote 7750 SR. This is the service label that the far end is expecting to receive for *service-id* when sending to the originating 7750 SR. If *service-id* does not exist in the remote 7750 SR, N/A is displayed. If *service-id* exists, but an ingress service label has not been assigned in the remote 7750 SR, Non-Existent is displayed.	*resp-ingress-vc-label* N/A Non-Existent
Responders Ingress Label Source	The assigned ingress service label source on the remote 7750 SR. If the ingress service label is manually defined on the remote 7750 SR, Manual is displayed. If the ingress service label is dynamically signaled on the remote 7750 SR, Signaled is displayed. If the *service-id* does not exist on the remote 7750 SR, N/A is displayed.	Manual Signaled N/A
Responders Ingress Service Label State	The assigned ingress service label state on the remote 7750 SR. If the remote 7750 SR considers its ingress service label operational, Up is displayed. If the remote 7750 SR considers its ingress service label inoperative, Down is displayed. If the *service-id* does not exist on the remote 7750 SR or the ingress service label has not been assigned on the remote 7750 SR, N/A is displayed.	Up Down N/A

Parameters *ip-address* — The far-end IP address to which to send the **svc-ping** request message in dotted decimal notation.

service *service-id* **—** The Service ID of the service being tested must be indicated with this parameter. The Service ID need not exist on the local 7750 SR to receive a reply message.

Values 1 — 2147483647

local-sdp — Specifies the **svc-ping** request message should be sent using the same service tunnel encapsulation labeling as service traffic. If **local-sdp** is specified, the command attempts to use an egress SDP-ID bound to the service with the specified **far-end** IP address with the VC-Label for the service. The far-end address of the specified SDP-ID is the expected *responder-id* within the reply received. The SDP-ID defines the encapsulation of the SDP tunnel encapsulation used to reach the far end; this can be IP/GRE or MPLS. On originator egress, the Service-ID must have an associated VC-Label to reach the far-end address of the SDP-ID and the SDP-ID must be operational for the message to be sent.

If **local-sdp** is not specified, the **svc-ping** request message is sent with GRE encapsulation with the OAM label.

The table below indicates whether a message is sent and how the message is encapsulated based on the state of the Service ID.

Local Service State	local-sdp Not Specified		local-sdp Specified	
	Message Sent	Message Encapsulation	Message Sent	Message Encapsulation
Invalid Local Service	Yes	Generic IP/GRE OAM (PLP)	No	None
No Valid SDP-ID Bound	Yes	Generic IP/GRE OAM (PLP)	No	None
SDP-ID Valid But Down	Yes	Generic IP/GRE OAM (PLP)	No	None
SDP-ID Valid and Up, But No Service Label	Yes	Generic IP/GRE OAM (PLP)	No	None
SDP-ID Valid, Up and Egress Service Label	Yes	Generic IP/GRE OAM (PLP)	Yes	SDP Encapsulation with Egress Service Label (SLP)

remote-sdp — Specifies **svc-ping** reply message from the **far-end** should be sent using the same service tunnel encapsulation labeling as service traffic.

If **remote-sdp** is specified, the **far-end** responder attempts to use an egress SDP-ID bound to the service with the message originator as the destination IP address with the VC-Label for the service. The SDP-ID defines the encapsulation of the SDP tunnel encapsulation used to reply to the originator; this can be IP/GRE or MPLS. On responder egress, the Service-ID must have an associated VC-Label to reach the originator address of the SDP-ID and the SDP-ID must be operational for the message to be sent.

If **remote-sdp** is not specified, the **svc-ping** request message is sent with GRE encapsulation with the OAM label.

The table below indicates how the message response is encapsulated based on the state of the remote Service ID.

Remote Service State	Message Encapsulation	
	remote-sdp Not Specified	remote-sdp Specified
Invalid Ingress Service Label	Generic IP/GRE OAM (PLP)	Generic IP/GRE OAM (PLP)
Invalid Service-ID	Generic IP/GRE OAM (PLP)	Generic IP/GRE OAM (PLP)
No Valid SDP-ID Bound on Service-ID	Generic IP/GRE OAM (PLP)	Generic IP/GRE OAM (PLP)
SDP-ID Valid But Down	Generic IP/GRE OAM (PLP)	Generic IP/GRE OAM (PLP)
SDP-ID Valid and Up, but No Service Label	Generic IP/GRE OAM (PLP)	Generic IP/GRE OAM (PLP)

Remote Service State	Message Encapsulation	
	remote-sdp Not Specified	**remote-sdp Specified**
SDP-ID Valid and Up, Egress Service Label, but VC-ID Mismatch	Generic IP/GRE OAM (PLP)	Generic IP/GRE OAM (PLP)
SDP-ID Valid and Up, Egress Service Label, but VC-ID Match	Generic IP/GRE OAM (PLP)	SDP Encapsulation with Egress Service Label (SLP)

Sample Output

```
*A:router1> svc-ping far-end 10.10.10.10 service 101 local-sdp remote-sdp
Request Result: Sent – Reply Received

Service-ID: 101

Err       Basic Info              Local     Remote
---       -----------------       ------    ------
__        Type:                   TLS       TLS
__        Admin State:            Up        Up
__        Oper State:             Up        Up
__        Service-MTU:            1514      1514
__        Customer ID:            1001      1001

Err       System IP Interface Info
---       ------------------------------------------------------------
Local Interface Name: “7750 SR-System-IP-Interface (Up to 32 chars)…”
__        Local IP Interface State:         Up
__        Local IP Address:                 10.10.10.11
__        IP Address Expected By Remote:    10.10.10.11
__        Expected Remote IP Address:       10.10.10.10
__        Actual Remote IP Address:         10.10.10.10

Err       SDP-ID Info             Local     Remote
---       -----------------       ------    ------
__        Path Used:              Yes       Yes
__        SDP-ID:                 123       325
__        Administrative State:   Up        Up
__        Operative State:        Up        Up
__        Binding Admin State:    Up        Up
__        Binding Oper State:     Up        Up
__        Binding VC-ID:          101       101

Err       Service Label Information   Label     Source        State
---       -------------------------   -----     -----------   -----
__        Local Egress Label:         45        Signaled      Up
__        Remote Expected Ingress:    45        Signaled      Up
__        Remote Egress:              34        Signaled      Up
__        Local Expected Ingress:     34        Signaled      Up
```

host-connectivity-verify

Syntax **host-connectivity-verify service** *service-id* [**sap** *sap-id*]

host-connectivity-verify subscriber *sub-ident-string* [**sla-profile** *sla-profile-name*]

Context <GLOBAL>

Description This command enables host connectivity verification checks.

Parameters **service** *service-id* **—** Specifies the service ID to diagnose or manage.

Values 1 — 2147483647

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

Values

sap-id:	null	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]	
	dot1q	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]:*qtag1*	
	qinq	[*port-id* | *bundle-id* | *lag-id*]:*qtag1.qtag2*	
	atm	[*port-id* | *bundle-id*][:*vpi/vci* | *vpi* | *vpi1.vpi2*]	
	frame	[*port-id* | *bundle-id*]:*dlci*	
	cisco-hdlc	*slot/mda/port.channel*	
	ima-grp	*bundle-id*[:*vpi/vci* | *vpi* | *vpi1.vpi2*]	
	port-id	*slot/mda/port*[.*channel*]	
	aps-id	aps-*group-id*[.*channel*]	
		aps	keyword
		group-id	1 — 64
	bundle-*type-slot/mda.bundle-num*		
		bundle	keyword
		type	ima, ppp
		bundle-num	1 — 128
	ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
		ccag	keyword
		id	1 — 8
		path-id	a, b
		cc-type	.sap-net, .net-sap]
		cc-id	0 — 4094
	lag-id	lag-*id*	
		lag	keyword
		id	1 — 64
	qtag1	0 — 4094	
	qtag2	*, 0 — 4094	
	vpi	NNI	0 — 4095
		UNI	0 — 255
	vci	1, 2, 5 — 65535	
	dlci	16 — 1022	

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 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

bundle-id — Specifies the multilink bundle to be associated with this IP interface. The **bundle** keyword must be entered at the beginning of the parameter.
The command syntax must be configured as follows:

bundle-id: **bundle-***type-slot-id*/*mda-slot*.*bundle-num*
bundle-id value range: 1 — 128

For example:

```
*A:ALA-12>config#  port bundle-ima-5/1.1
*A:ALA-12>config>port# multilink-bundle
```

ima — Specifies Inverse Multiplexing over ATM. An IMA group is a collection of physical links bundled together and assigned to an ATM port.

qtag1, qtag2 — Specifies the encapsulation value used to identify the SAP on the port or sub-port. If this parameter is not specificially defined, the default value is 0.

Values qtag1: 0 — 4094
qtag2 : * | 0 — 4094

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 TDM	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 TDM	BCP-Dot1q	0 — 4094	The SAP is identified by the 802.1Q tag on the channel.
SONET/SDH TDM	Frame Relay	16 — 991	The SAP is identified by the data link connection identifier (DLCI).
SONET/SDH ATM	ATM	vpi (NNI) 0 — 4095 vpi (UNI) 0 — 255 vci 1, 2, 5 — 65535	The SAP is identified by the PVC identifier (vpi/vci).

sub-profile *sub-profile-name* **—** Specifies an existing subscriber profile name. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

sla-profile *sla-profile-name* **—** Specifies an existing SLA profile name. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

vprn-ping

Syntax **vprn-ping** *service-id* **source** *ip-address* **destination** *ip-address* [**fc** *fc-name* [**profile** [**in** | **out**]] [**size** *size*] [**ttl** *vc-label-ttl*] [**return-control**] [**interval** *interval*] [**count** *send-count*] [**timeout** *timeout*]

Context <GLOBAL>
config>saa>test>type

Description This command performs a VPRN ping.

Parameters **service** *service-id* **—** The VPRN service ID to diagnose or manage.

Values 1 — 2147483647

source *ip-address* **—** The IP prefix for the source IP address in dotted decimal notation.

Values 0.0.0.0 — 255.255.255.255

destination *ip-address* **—** The IP prefix for the destination IP address in dotted decimal notation.

Values 0.0.0.0 — 255.255.255.255

fc-name — The forwarding class of the MPLS echo request encapsulation.

Default be

Values be, l2, af, l1, h2, ef, h1, nc

profile {**in** | **out**} **—** The profile state of the MPLS echo request encapsulation.

Default out

size *octets* **—** The OAM request packet size in octets, expressed as a decimal integer.

Values 1 — 65535

ttl *vc-label-ttl* **—** The TTL value in the VC label for the OAM request, expressed as a decimal integer.

Default 255

Values 1 — 255

return-control — Specifies the response to come on the control plane.

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second where the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

count *send-count* **—** The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

Sample Output

```
A:PE_1# oam vprn-ping 25 source 10.4.128.1 destination 10.16.128.0
Sequence Node-id                                Reply-Path Size     RTT
-------------------------------------------------------------------------------
[Send request Seq. 1.]
1        10.128.0.3:cpm                          In-Band   100      0ms
-------------------------------------------------------------------------------
...
A:PE_1#
-------------------------------------------------------------------------------
```

vprn-trace

Syntax **vprn-trace** *service-id* **source** *src-ip* **destination** *ip-address* [**fc** *fc-name* [**profile** [**in** | **out**]] [**size** *size*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**return-control**] [**probe-count** *probes-per-hop*] [**interval** *seconds*] [**timeout** *timeout*]

Context <GLOBAL>
config>saa>test>type

Description Performs VPRN trace.

Parameters **service** *service-id* **—** The VPRN service ID to diagnose or manage.

Values 1 — 2147483647

source *src-ip* **—** The IP prefix for the source IP address in dotted decimal notation.

Values 0.0.0.0 — 255.255.255.255

fc-name — The forwarding class of the MPLS echo request encapsulation.

Default be

Values be, l2, af, l1, h2, ef, h1, nc

profile {**in** | **out**} **—** The profile state of the MPLS echo request encapsulation.

Default out

destination *dst-ip* **—** The IP prefix for the destination IP address in dotted decimal notation.

Values 0.0.0.0 — 255.255.255.255

size *octets* **—** The OAM request packet size in octets, expressed as a decimal integer.

min-ttl *vc-label-ttl* **—** The minimum TTL value in the VC label for the trace test, expressed as a decimal integer.

Default 1

Values 1 — 255

max-ttl *vc-label-ttl* **—** The maximum TTL value in the VC label for the trace test, expressed as a decimal integer.

Default 4

Values 1 — 255

return-control — Specifies the OAM reply to a data plane OAM request be sent using the control plane instead of the data plane.

Default OAM reply sent using the data plane.

probe-count *send-count* **—** The number of OAM requests sent for a particular TTL value, expressed as a decimal integer.

Default 1

Values 1 — 10

interval *seconds* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second where the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 3

Values 1 — 10

Sample Output

```
A:PE_1# oam vprn-trace 25 source 10.4.128.1 destination 10.16.128.0
TTL Seq Reply Node-id         Rcvd-on         Reply-Path RTT
-------------------------------------------------------------------------------
[Send request TTL: 1, Seq. 1.]
1   1   1     10.128.0.4      cpm             In-Band    0ms
  Requestor 10.128.0.1 Route: 0.0.0.0/0
    Vpn Label: 131071 Metrics 0 Pref 170 Owner bgpVpn
    Next Hops: [1] ldp tunnel
    Route Targets: [1]: target:65100:1
  Responder 10.128.0.4 Route: 10.16.128.0/24
    Vpn Label: 131071 Metrics 0 Pref 170 Owner bgpVpn
    Next Hops: [1] ldp tunnel
    Route Targets: [1]: target:65001:100

[Send request TTL: 2, Seq. 1.]
2   1   1     10.128.0.3      cpm             In-Band    0ms
  Requestor 10.128.0.1 Route: 0.0.0.0/0
    Vpn Label: 131071 Metrics 0 Pref 170 Owner bgpVpn
    Next Hops: [1] ldp tunnel
    Route Targets: [1]: target:65100:1
  Responder 10.128.0.3 Route: 10.16.128.0/24
    Vpn Label: 0 Metrics 0 Pref 0 Owner local
    Next Hops: [1] ifIdx 2 nextHopIp 10.16.128.0

[Send request TTL: 3, Seq. 1.]
[Send request TTL: 4, Seq. 1.]
...
-------------------------------------------------------------------------------
A:PE_1#
```

VPLS MAC Diagnostics

mac-populate

Syntax **mac-populate** *service-id* **mac** *ieee-address* [**flood**] [**age** *seconds*] [**force**]

Context oam

Description Populates the FIB with an OAM-type MAC entry indicating the node is the egress node for the MAC address and optionally floods the OAM MAC association throughout the service.

A **mac-populate** installs an OAM MAC into the service FIB indicating the device is the egress node for a particular MAC address. The MAC address can be bound to a particular SAP (the **target-sap**) or can be associated with the control plane in that any data destined to the MAC address is forwarded to the control plane (cpm). As a result, if the service on the node has neither a FIB nor an egress SAP, then it is not allowed to initiate a **mac-populate**.

The MAC address that is populated in the FIBs in the provider network is given a type OAM, so that it can be treated distinctly from regular dynamically learned or statically configured MACs. Note that OAM MAC addresses are operational MAC addresses and are not saved in the device configuration. An exec file can be used to define OAM MACs after system initialization.

The **force** option in **mac-populate** forces the MAC in the table to be type OAM in the case it already exists as a dynamic, static or an OAM induced learned MAC with some other type binding.

An OAM-type MAC cannot be overwritten by dynamic learning and allows customer packets with the MAC to either ingress or egress the network while still using the OAM MAC entry.

The **flood** option causes each upstream node to learn the MAC (that is, populate the local FIB with an OAM MAC entry) and to flood the request along the data plane using the flooding domain.The flooded **mac-populate** request can be sent via the data plane or the control plane. The **send-control** option specifies the request be sent using the control plane. If **send-control** is not specified, the request is sent using the data plane.

An **age** can be provided to age a particular OAM MAC using a specific interval. By default, OAM MAC addresses are not aged and can be removed with a **mac-purge** or with an FDB clear operation.

When split horizon group (SHG) is configured, the flooding domain depends on which SHG the packet originates from. The **target-sap** *sap-id* value dictates the originating SHG information.

Parameters **service** *service-id* **—** The Service ID of the service to diagnose or manage.

Values 1 — 2147483647

destination *ieee-address* **—** The MAC address to be populated.

flood — Sends the OAM MAC populate to all upstream nodes.

Default MAC populate only the local FIB.

age *seconds* **—** The age for the OAM MAC, expressed as a decimal integer.

Default The OAM MAC does not age.

Values 1 — 65535

force — Converts the MAC to an OAM MAC even if it currently another type of MAC.

Default Do not overwrite type.

target-sap *sap-id* **—** The local target SAP bound to a service on which to associate the OAM MAC. By default, the OAM MAC is associated with the control place, that is, it is associated with the CPU on the router.

When the **target-sap** *sap-id* value is not specified the MAC is bound to the CPM. The originating SHG is o (zero). When the **target-sap** *sap-id* value is specified, the originating SHG is the SHG of the target-sap.

Default Associate OAM MAC with the control plane (cpu)

Values

sap-id:	null	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]	
	dot1q	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]:*qtag1*	
	qinq	[*port-id* | *bundle-id* | *lag-id*]:*qtag1.qtag2*	
	atm	[*port-id* | *bundle-id*][:*vpi/vci* | *vpi* | *vpi1.vpi2*]	
	frame	[*port-id* | *bundle-id*]:*dlci*	
	cisco-hdlc	*slot/mda/port.channel*	
	ima-grp	*bundle-id*[:*vpi/vci* | *vpi* | *vpi1.vpi2*]	
	port-id	*slot/mda/port*[.*channel*]	
	aps-id	aps-*group-id*[.*channel*]	
		aps	keyword
		group-id	1 — 64
	bundle-*type-slot/mda.bundle-num*		
		bundle	keyword
		type	ima, ppp
		bundle-num	1 — 128
	ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
		ccag	keyword
		id	1 — 8
		path-id	a, b
		cc-type	.sap-net, .net-sap]
		cc-id	0 — 4094
	lag-id	lag-*id*	
		lag	keyword
		id	1 — 64
	qtag1	0 — 4094	
	qtag2	*, 0 — 4094	
	vpi	NNI	0 — 4095
		UNI	0 — 255
	vci	1, 2, 5 — 65535	
	dlci	16 — 1022	

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 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

bundle-id — Specifies the multilink bundle to be associated with this IP interface. The **bundle** keyword must be entered at the beginning of the parameter.
The command syntax must be configured as follows:

bundle-id: **bundle**-*type-slot-id*/*mda-slot*.*bundle-num*
bundle-id value range: 1 — 128

For example:

```
*A:ALA-12>config#  port bundle-ima-5/1.1
*A:ALA-12>config>port# multilink-bundle
```

ima — Specifies Inverse Multiplexing over ATM. An IMA group is a collection of physical links bundled together and assigned to an ATM port.

qtag1, qtag2 — Specifies the encapsulation value used to identify the SAP on the port or sub-port. If this parameter is not specificially defined, the default value is 0.

Values qtag1: 0 — 4094
qtag2 : * | 0 — 4094

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 TDM	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 TDM	BCP-Dot1q	0 — 4094	The SAP is identified by the 802.1Q tag on the channel.
SONET/SDH TDM	Frame Relay	16 — 991	The SAP is identified by the data link connection identifier (DLCI).
SONET/SDH ATM	ATM	vpi (NNI) 0 — 4095 vpi (UNI) 0 — 255 vci 1, 2, 5 — 65535	The SAP is identified by the PVC identifier (vpi/vci).

mac-purge

Syntax **mac-purge** *service-id* **target** *ieee-address* [**flood**] [**send-control**] [**register**]

Context <GLOBAL>
config>saa>test>type

Description Removes an OAM-type MAC entry from the FIB and optionally floods the OAM MAC removal throughout the service.

A **mac-purge** can be sent via the forwarding path or via the control plane. When sending the MAC purge using the data plane, the TTL in the VC label is set to 1. When sending the MAC purge using the control plane, the packet is sent directly to the system IP address of the next hop.

A MAC address is purged only if it is marked as OAM.

A mac-purge request is an HVPLS OAM packet, with the following fields. The Reply Flags is set to 0 (since no reply is expected), the Reply Mode and Reserved fields are set to 0. The Ethernet header has source set to the (system) MAC address, the destination set to the broadcast MAC address. There is a VPN TLV in the FEC Stack TLV to identify the service domain.

If the register option is provided, the R bit in the Address Delete flags is turned on.

The **flood** option causes each upstream node to be sent the OAM MAC delete request and to flood the request along the data plane using the flooding domain. The flooded **mac-purge** request can be sent via the data plane or the control plane. The **send-control** option specifies the request be sent using the control plane. If **send-control** is not specified, the request is sent using the data plane.

The **register** option reserves the MAC for OAM testing where it is no longer an active MAC in the FIB for forwarding, but it is retained in the FIB as a registered OAM MAC. Registering an OAM MAC prevents relearns for the MAC based on customer packets. Relearning a registered MAC can only be done through a **mac-populate** request. The originating SHG is always 0 (zero).

Parameters **service** *service-id* **—** The Service ID of the service to diagnose or manage.

Values 1 — 2147483647

target *ieee-address* **—** The MAC address to be purged.

flood — Sends the OAM MAC purge to all upstream nodes.

Default MAC purge only the local FIB.

send-control — Send the mac-purge request using the control plane.

Default Request is sent using the data plane.

register — Reserve the MAC for OAM testing.

Default Do not register OAM MAC.

IGMP Snooping Diagnostics

mfib-ping

Syntax **mfib-ping service** *service-id* **source** *src-ip* **destination** *mcast-address* [**size** *size*] [**ttl** *vc-label-ttl*] [**return-control**] [**interval** *interval*] [**count** *send-count*] [**timeout** *timeout*]

Context oam
config>saa>test>type

Description The mfib-ping utility determines the list of SAPs which egress a certain IP multicast stream (identified by source unicast and destination multicast IP addresses) within a VPLS service.

An mfib-ping packet is always sent via the data plane.

An mfib-ping is forwarded across the VPLS following the MFIB. If an entry for the specified source unicast and destination multicast IP addresses exist in the MFIB for that VPLS, then the packet is forwarded along those paths, provided they are active. A response is generated only when there is an egress SAP binding for the specified IP multicast stream.

An mfib-ping reply can be sent using the data plane or the control plane. The return-control option specifies the reply be sent using the control plane. If return-control is not specified, the reply is sent using the data plane.

Parameters **service** *service-id* **—** The service ID of the VPLS to diagnose or manage.

Values 1 — 2147483647

source *src-ip* **—** The source IP address for the OAM request.

destination *mcast-address* **—** The destination multicast address for the OAM request.

size *size* **—** The multicast OAM request packet size in octets, expressed as a decimal integer. The request payload is padded to the specified size with a 6 byte PAD header and a byte payload of 0xAA as necessary. If the octet size specified is less than the minimum packet, the minimum sized packet necessary to send the request is used.

Default No OAM packet padding.

Values 1 — 65535

ttl *vc-label-ttl* **—** The TTL value in the VC label for the OAM request, expressed as a decimal integer.

Default 255

Values 1 — 255

return-control **—** Specifies the OAM reply has to be sent using the control plane instead of the data plane.

Default OAM reply is sent using the data plane.

interval *interval* **—** The interval parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the interval is set to 1 second where the timeout value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

count *send-count* **—** The number of messages to send, expressed as a decimal integer. The count parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message interval value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *seconds* **—** The timeout parameter in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the 7750 SR will wait for a message reply after sending the next message request. Upon the expiration of message timeout, the requesting 7750 SR assumes that the message response will not be received. A 'request timeout' message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 100

Special Cases **MFIB 224.0.0.X pings —** Mfib-ping requests directed to a destination address in the special 224.0.0.X range are flooded throughout the service flooding domain and will receive a response from all operational SAPs. Note that SAPs that are operationally down do not reply. If EMG is enabled, mfib-ping will return only the first SAP in each chain.

Multicast FIB Connectivity Test Sample Output

```
A:ALA-A# oam mfib-ping service 10 source 10.10.10.1 destination 225.0.0.1 count 2
Seq Node-id                                                 Path     Size  RTT
-------------------------------------------------------------------------------
[Send request Seq. 1.]
1   51.51.51.51:sap1/1/1                                   Self     100   0ms
1   54.54.54.54:sap1/1/2                                   In-Band  100   20ms
1   54.54.54.54:sap1/1/3                                   In-Band  100   10ms
1   52.52.52.52:sap1/1/2                                   In-Band  100   10ms
1   52.52.52.52:sap1/1/3                                   In-Band  100   20ms
[Send request Seq. 2.]
2   51.51.51.51:sap1/1/1                                   Self     100   0ms
2   52.52.52.52:sap1/1/2                                   In-Band  100   10ms
2   54.54.54.54:sap1/1/2                                   In-Band  100   10ms
2   52.52.52.52:sap1/1/3                                   In-Band  100   20ms
2   54.54.54.54:sap1/1/3                                   In-Band  100   30ms
-------------------------------------------------------------------------------
A:ALA-AIM# oam mfib-ping service 1 source 11.11.0.0 destination 224.0.0.1
Seq Node-id                                                 Path     Size  RTT
-------------------------------------------------------------------------------
[Send request Seq. 1.]
1   10.20.1.3:sap1/1/5:1                       Not in MFIB Self     40    0ms
1   10.20.1.3:sap1/1/2:1                                   Self     40    10ms
[Echo replies received: 2]
-------------------------------------------------------------------------------
```

```
A:ALA-AIM#
```

EFM Commands

efm

Syntax **efm** *port-id*

Context oam

Description This command enables Ethernet in the First Mile (EFM) OAM tests loopback tests on the specified port. The EFM OAM remote loopback OAMPDU will be sent to the peering device to trigger remote loopback.

Parameters *port-id* — Specify the port ID in the slot/mda/port format.

local-loopback

Syntax **local-loopback** {**start** | **stop**}

Context oam>emf

Description This command enables local loopback tests on the specified port.

remote-loopback

Syntax **remote-loopback** {**start** | **stop**}

Context oam>emf

Description This command enables remote Ethernet in the First Mile (EFM) OAM loopback tests on the specified port. The EFM OAM remote loopback OAMPDU will be sent to the peering device to trigger remote loopback.

In order for EFM OAM tunneling to function properly, EFM OAM tunneling should be configured for VLL services or a VPLS service with two SAPs only.

Service Assurance Agent (SAA) Commands

saa

Syntax **saa**

Context config

Description This command creates the context to configure the Service Assurance Agent (SAA) tests.

test

Syntax **test** *name* [**owner** *test-owner*]
no test *name*

Context config>saa

Description This command identifies a test and create/modify the context to provide the test parameters for the named test. Subsequent to the creation of the test instance the test can be started in the OAM context.

A test can only be modified while it is shut down.

The **no** form of this command removes the test from the configuration. In order to remove a test it can not be active at the time.

Parameters *name* — Identify the saa test name to be created or edited.

owner *test-owner* **—** Specifies the owner of an SAA operation up to 32 characters in length.

Values If a *test-owner* value is not specified, tests created by the CLI have a default owner "TiMOS CLI".

description

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

Context config>saa>test

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

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

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

Default No description associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters 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.

jitter-event

Syntax **jitter-event rising-threshold** *threshold* [**falling-threshold** *threshold*] [*direction*]
no jitter-event

Context config>saa>test

Description Specifies that at the termination of an SAA test probe, the calculated jitter value is evaluated against the configured rising and falling jitter thresholds. SAA threshold events are generated as required.

The configuration of jitter event thresholds is optional.

Parameters **rising-threshold** *threshold* **—** Specifies a rising threshold jitter value. When the test run is completed, the calculated jitter value is compared to the configured jitter rising threshold. If the test run jitter value is greater than the configured rising threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483647 milliseconds

falling-threshold *threshold* **—** Specifies a falling threshold jitter value. When the test run is completed, the calculated jitter value is compared to the configured jitter falling threshold. If the test run jitter value is greater than the configured falling threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483647 milliseconds

direction — Specifies the direction for OAM ping responses received for an OAM ping test run.

Values **inbound** — Monitor the value of jitter calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run.
outbound — Monitor the value of jitter calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run.
roundtrip — Monitor the value of jitter calculated for the round trip, two-way, OAM ping requests and replies for an OAM ping test run.

Default roundtrip

latency-event

Syntax **latency-event rising-threshold** *threshold* [**falling-threshold** *threshold*] [*direction*]
no latency-event

Context config>saa>test

Description Specifies that at the termination of an SAA test probe, the calculated latency event value is evaluated against the configured rising and falling latency event thresholds. SAA threshold events are generated as required.

The configuration of latency event thresholds is optional.

Parameters **rising-threshold** *threshold* **—** Specifies a rising threshold latency value. When the test run is completed, the calculated latency value is compared to the configured latency rising threshold. If the test run latency value is greater than the configured rising threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483647 milliseconds

falling-threshold *threshold* **—** Specifies a falling threshold latency value. When the test run is completed, the calculated latency value is compared to the configured latency falling threshold. If the test run latency value is greater than the configured falling threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483647 milliseconds

direction — Specifies the direction for OAM ping responses received for an OAM ping test run.

Values **inbound** — Monitor the value of jitter calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run.
outbound — Monitor the value of jitter calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run.
roundtrip — Monitor the value of jitter calculated for the round trip, two-way, OAM ping requests and replies for an OAM ping test run.

Default roundtrip

loss-event

Syntax **loss-event rising-threshold** *threshold* [**falling-threshold** *threshold*] [*direction*]
no loss-event

Context config>saa>test

Description Specifies that at the termination of an SAA testrun, the calculated loss event value is evaluated against the configured rising and falling loss event thresholds. SAA threshold events are generated as required.

The configuration of loss event thresholds is optional.

Parameters **rising-threshold** *threshold* **—** Specifies a rising threshold loss event value. When the test run is completed, the calculated loss event value is compared to the configured loss event rising threshold. If the test run loss event value is greater than the configured rising threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483647 packets

falling-threshold *threshold* **—** Specifies a falling threshold loss event value. When the test run is completed, the calculated loss event value is compared to the configured loss event falling threshold. If the test run loss event value is greater than the configured falling threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483647 packets

direction — Specifies the direction for OAM ping responses received for an OAM ping test run.

Values **inbound** — Monitor the value of jitter calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run.
outbound — Monitor the value of jitter calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run.
roundtrip — Monitor the value of jitter calculated for the round trip, two-way, OAM ping requests and replies for an OAM ping test run.

Default roundtrip

type

Syntax **type**
no type

Context config>saa>test

Description This command creates the context to provide the test type for the named test. Only a single test type can be configured.

A test can only be modified while the test is in shut down mode.

Once a test type has been configured the command can be modified by re-entering the command, the test type must be the same as the previously entered test type.

To change the test type, the old command must be removed using the **config>saa>test>no type** command.

cpe-ping

Syntax **cpe-ping service** *service-id* **destination** *ip-address* **source** *ip-address* [**ttl** *vc-label-ttl*] [**return-control**] [**source-mac** *ieee-address*] [**fc** *fc-name* [**profile** [**in** | **out**]] [**interval** *interval*] [**count** *send-count*] [**send-control]**

Context <GLOBAL>
config>saa>test>type

Description This ping utility determines the IP connectivity to a CPE within a specified VPLS service.

Parameters **service** *service-id* **—** The service ID of the service to diagnose or manage.

Values 1 — 2147483647

destination *ip-address* **—** Specifies the IP address to be used as the destination for performing an OAM ping operations.

source *ip-address* **—** Specify an unused IP address in the same network that is associated with the VPLS.

ttl *vc-label-ttl* **—** The TTL value in the VC label for the OAM MAC request, expressed as a decimal integer.

Default 255

Values 1 — 255

return-control — Specifies the MAC OAM reply to a data plane MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM reply sent using the data plane.

source-mac *ieee-address* **—** Specify the source MAC address that will be sent to the CPE. If not specified or set to 0, the MAC address configured for the CPM is used.

fc-name — The forwarding class of the MPLS echo request encapsulation.

Default be

Values be, l2, af, l1, h2, ef, h1, nc

profile {**in** | **out**} **—** The profile state of the MPLS echo request encapsulation.

Default out

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second where the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

count *send-count* **—** The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 255

send-control — Specifies the MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM request sent using the data plane.

dns

Syntax **dns target-addr** *dns-name* **name-server** *ip-address* [**source** *ip-address*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]

Context <GLOBAL>
config>saa>test>type

Description This command configures a DNS name resolution test.

Parameters **target-addr —** The IP host address to be used as the destination for performing an OAM ping operation.

dns-name — The DNS name to be resolved to an IP address.

name-server *ip-address* **—** Specifies the server connected to a network that resolves network names into network addresses.

source *ip-address* **—** Specifies the IP address to be used as the source for performing an OAM ping operation.

send-count *send-count* **—** The number of messages to send, expressed as a decimal integer. The **send-count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

icmp-ping

Syntax **icmp-ping** [*ip-address* | *dns-name*] [**rapid** | **detail**] [**ttl** *time-to-live*] [**tos** *type-of-service*] [**size** *bytes*] [**pattern** *pattern*] [**source** *ip-address* | *dns-name*] [**interval** *seconds*] [{**next-hop** *ip-*

address} | {**interface** *interface-name*} | **bypass-routing**] [**count** *requests*] [**do-not-fragment**] [**router** *router-instance*] [**timeout** *timeout*]

Context config>saa>test>type

Description This command configures an ICMP traceroute test.

Parameters *ip-address* — The far-end IP address to which to send the **svc-ping** request message in dotted decimal notation.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x
	x:x:x:x:x:x:d.d.d.d
	x: [0 — FFFF]H
	d: [0 — 255]D

dns-name — The DNS name of the far-end device to which to send the **svc-ping** request message, expressed as a character string up to 63 characters maximum.

Values

ipv6-address:	x: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, mandatory for link local addresses)

rapid — Packets will be generated as fast as possible instead of the default 1 per second.

detail — Displays detailed information.

ttl *time-to-live* **—** The TTL value for the MPLS label, expressed as a decimal integer.

Values 1 — 128

tos *type-of-service* **—** Specifies the service type.

Values 0 — 255

size *bytes* **—** The request packet size in bytes, expressed as a decimal integer.

Values 0 — 16384

pattern ***pattern*** **—** The date portion in a ping packet will be filled with the pattern value specified. If not specified, position info will be filled instead.

Values 0 — 65535

source *ip-address*|*dns-name* **—** Specifies the IP address to be used.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x
	x:x:x:x:x:x:d.d.d.d
	x: [0 — FFFF]H
	d: [0 — 255]Ddns-name:128 characters max

interval *seconds* **—** This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

next-hop *ip-address* **—** Only displays static routes with the specified next hop IP address.

Values

ipv4-address:	a.b.c.d (host bits must be 0)
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 — FFFF]H
	d: [0 — 255]D

interface *interface-name* — The name used to refer to the interface. The name must already exist in the **config>router>interface** context.

bypass-routing — Specifies whether to send the ping request to a host on a directly attached network bypassing the routing table.

count *requests* **—** Specifies the number of times to perform an OAM ping probe operation. Each OAM echo message request must either timeout or receive a reply before the next message request is sent.

Values 1 — 100000

Default 5

do-not-fragment — Sets the DF (Do Not Fragment) bit in the ICMP ping packet.

router *router-instance* **—** Specifies the router name or service ID.

Values

router-name:	Base , management
service-id:	1 — 2147483647

Default Base

timeout *timeout* **—** Overrides the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

icmp-trace

Syntax **icmp-trace** [*ip-address* | *dns-name*] [**ttl** *time-to-live*] [**wait** *milli-seconds*] [**tos** *type-of-service*] [**source** *ip-address*] [**tos** *type-of-service*] [**router** *router-instance*]

Context config>saa>test>type

Description This command configures an ICMP traceroute test.

Parameters *ip-address* — The far-end IP address to which to send the **svc-ping** request message in dotted decimal notation.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x
	x:x:x:x:x:x:d.d.d.d
x:	[0 — FFFF]H
d:	[0 — 255]D

dns-name — The DNS name of the far-end device to which to send the **svc-ping** request message, expressed as a character string to 63 characters maximum.

ttl *time-to-live* **—** The TTL value for the MPLS label, expressed as a decimal integer.

Values 1 — 255

wait *milliseconds* **—** The time in milliseconds to wait for a response to a probe, expressed as a decimal integer.

Default 5000

Values 1 — 60000

tos *type-of-service* **—** Specifies the service type.

Values 0 — 255

source *ip-address* **—** Specifies the IP address to be used.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x
	x:x:x:x:x:x:d.d.d.d
x:	[0 — FFFF]H
d:	[0 — 255]D

router *router-instance* **—** Specifies the router name or service ID.

Values

router-name:	Base , management
service-id:	1 — 2147483647

Default Base

lsp-ping

Syntax **lsp-ping** {{[*lsp-name*] [**path** *path-name*]} | {**prefix** *ip-prefix/mask*}} [**fc** *fc-name*] [**profile** {**in** | **out**}]] [**size** *octets*] [**ttl** *label-ttl*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**path-destination** *ip-address*] [**interface** *if-name* | **next-hop** *ip-address*]][**detail**]

Context oam>
config>saa>test>type

Description Performs in-band LSP connectivity tests.

The **lsp-ping** command performs an LSP ping using the protocol and data structures defined in the RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*.

The LSP ping operation is modeled after the IP ping utility which uses ICMP echo request and reply packets to determine IP connectivity.

In an LSP ping, the originating device creates an MPLS echo request packet for the LSP and path to be tested. The MPLS echo request packet is sent through the data plane and awaits an MPLS echo reply packet from the device terminating the LSP. The status of the LSP is displayed when the MPLS echo reply packet is received.

Parameters *lsp-name* — Name that identifies an LSP to ping. The LSP name can be up to 32 characters long.

path *path-name* **—** The LSP pathname along which to send the LSP ping request.

Default The active LSP path.

Values Any path name associated with the LSP.

prefix *ip-prefix/mask* **—** Specifies the address prefix and subnet mask of the destination node.

fc *fc-name* **—** The **fc** parameter is used to indicate the forwarding class of the MPLS echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

The LSP-EXP mappings on the receive network interface controls the mapping back to the internal forwarding class used by the far-end 7750 SR that receives the message request. The egress mappings of the egress network interface on the far-end 7750 SR controls the forwarding class markings on the return reply message.

The LSP-EXP mappings on the receive network interface controls the mapping of the message reply back at the originating 7750 SR.

Default be

Values be, l2, af, l1, h2, ef, h1, nc

profile {**in** | **out**} **—** The profile state of the MPLS echo request encapsulation.

Default out

size *octets* **—** The MPLS echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeroes to the specified size.

Default 68 — The system sends the minimum packet size, depending on the type of LSP. No padding is added.

Values 84 — 65535

ttl *label-ttl* **—** The TTL value for the MPLS label, expressed as a decimal integer.

Default 255

Values 1 — 255

send-count *send-count* **—** The number of messages to send, expressed as a decimal integer. The **send-count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the

requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

path-destination *ip-address* **—**

interface *interface-name* — Specifies the name of an IP interface. The name must already exist in the **config**>**router**>**interface** context.

next-hop *ip-address* **—** Only displays static routes with the specified next hop IP address.

Values		
	ipv4-address:	a.b.c.d (host bits must be 0)
	ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x: [0 — FFFF]H
		d: [0 — 255]D

lsp-trace

Syntax **lsp-trace** {{[*lsp-name]* [**path** path-name]} **|** {**prefix** ip-prefix/mask}} [**fc** *fc-name*] [**profile** {**in** | **out**}]] [**max-fail** *no-response-count*] [**probe-count** *probes-per-hop*] [**size** *octets*][**min-ttl** *min-label-ttl*]] [**max-ttl** *max-label-ttl*] [**timeout** *timeout*] [[**interval** *interval*] [**path-destination** *ip-address*] [**interface** *if-name* | **next-hop** *ip-address*]][**detail**]

Context <GLOBAL>
config>saa>test>type

Description Displays the hop-by-hop path for an LSP.

The **lsp-trace** command performs an LSP traceroute using the protocol and data structures defined in the IETF draft (draft-ietf-mpls-lsp-ping-02.txt).

The LSP traceroute operation is modeled after the IP traceroute utility which uses ICMP echo request and reply packets with increasing TTL values to determine the hop-by-hop route to a destination IP.

In an LSP traceroute, the originating device creates an MPLS echo request packet for the LSP to be tested with increasing values of the TTL in the outermost label. The MPLS echo request packet is sent through the data plane and awaits a TTL exceeded response or the MPLS echo reply packet from the device terminating the LSP. The devices that reply to the MPLS echo request packets with the TTL exceeded and the MPLS echo reply are displayed.

Parameters *lsp-name* — Name that identifies an LSP to ping. The LSP name can be up to 32 characters long.

path *path-name* **—** The LSP pathname along which to send the LSP trace request.

Default The active LSP path.

Values Any path name associated with the LSP.

prefix *ip-prefix/mask* **—** Specifies the address prefix and subnet mask of the destination node.

size *octets* **—** The MPLS echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeroes to the specified size.

Default 68 — The system sends the minimum packet size, depending on the type of LSP. No padding is added.

Values 84 — 65535

min-ttl *min-label-ttl* **—** The minimum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer.

Default 1

Values 1 — 255

max-ttl *max-label-ttl* **—** The maximum TTL value in the MPLS label for the LDP treetrace test, expressed as a decimal integer.

Default 30

Values 1 — 255

max-fail *no-response-count* **—** The maximum number of consecutive MPLS echo requests, expressed as a decimal integer that do not receive a reply before the trace operation fails for a given TTL.

Default 5

Values 1 — 255

send-count *send-count* **—** The number of messages to send, expressed as a decimal integer. The **send-count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the 7750 SR will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting 7750 SR assumes that the message response will not be received. A 'request timeout' message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 3

Values 1 — 10

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

fc *fc-name* **—** The **fc** parameter is used to indicate the forwarding class of the MPLS echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

The LSP-EXP mappings on the receive network interface controls the mapping back to the internal forwarding class used by the far-end 7750 SR that receives the message request. The egress mappings of the egress network interface on the far-end 7750 SR controls the forwarding class markings on the return reply message.

The LSP-EXP mappings on the receive network interface controls the mapping of the message reply back at the originating 7750 SR.

fc-name — The forwarding class of the MPLS echo request encapsulation.

Default be

Values be, l2, af, l1, h2, ef, h1, nc

profile {**in** | **out**} **—** The profile state of the MPLS echo request encapsulation.

Default out

mac-ping

Syntax **mac-ping service** *service-id* **destination** *dst-ieee-address* [**source** *src-ieee-address*] [**fc** *fc-name* [**profile in** | **out**]] [**size** *octets*] [**ttl** *vc-label-ttl*] [**count** *send-count*] [**send-control**] [**return-control**] [**interval** *interval*] [**timeout** *timeout*]

Context oam
config>saa>test>type

Description The mac-ping utility is used to determine the existence of an egress SAP binding of a given MAC within a VPLS service.

A **mac-ping** packet can be sent via the control plane or the data plane. The **send-control** option specifies the request be sent using the control plane. If **send-control** is not specified, the request is sent using the data plane.

A **mac-ping** is forwarded along the flooding domain if no MAC address bindings exist. If MAC address bindings exist, then the packet is forwarded along those paths, provided they are active. A response is generated only when there is an egress SAP binding for that MAC address or if the MAC address is a “local” OAM MAC address associated with the device’s control plan.

A **mac-ping** reply can be sent using the data plane or the control plane. The **return-control** option specifies the reply be sent using the control plane. If **return-control** is not specified, the request is sent using the data plane.

A **mac-ping** with data plane reply can only be initiated on nodes that can have an egress MAC address binding. A node without a FIB and without any SAPs cannot have an egress MAC address binding, so it is not a node where replies in the data plane will be trapped and sent up to the control plane.

A control plane request is responded to via a control plane reply only.

By default, MAC OAM requests are sent with the system or chassis MAC address as the source MAC. The **source** option allows overriding of the default source MAC for the request with a specific MAC address.

When a **source** *ieee-address* value is specified and the source MAC address is locally registered within a split horizon group (SHG), then this SHG membership will be used as if the packet originated from this SHG. In all other cases, SHG 0 (zero) will be used. Note that if the **mac-trace** is originated from a non-zero SHG, such packets will not go out to the same SHG.

If EMG is enabled, mac-ping will return only the first SAP in each chain.

Parameters **service** *service-id* **—** The service ID of the service to diagnose or manage.

Values 1 — 2147483647

destination *ieee-address* **—** The destination MAC address for the OAM MAC request.

size *octets* **—** The MAC OAM request packet size in octets, expressed as a decimal integer. The request payload is padded to the specified size with a 6 byte PAD header and a byte payload of 0xAA as necessary. If the octet size specified is less than the minimum packet, the minimum sized packet necessary to send the request is used.

Default No OAM packet padding.

Values 1 — 65535

ttl *vc-label-ttl* **—** The TTL value in the VC label for the OAM MAC request, expressed as a decimal integer.

Default 255

Values 1 — 255

send-control — Specifies the MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM request sent using the data plane.

return-control — Specifies the MAC OAM reply to a data plane MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM reply sent using the data plane.

source *src-ieee-address* **—** The source MAC address from which the OAM MAC request originates. By default, the system MAC address for the chassis is used.

Default The system MAC address.

Values Any unicast MAC value.

fc *fc-name* **—** The **fc** parameter is used to test the forwarding class of the MPLS echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

Values be, l2, af, l1, h2, ef, h1, nc

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second where the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

count *send-count* **—** The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

mac-trace

Syntax **mac-trace service** *service-id* **destination** *ieee-address* [**size** *octets*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**send-control**] [**return-control**] [**source** *ieee-address*] [**probe-count** *probes-per-hop*] [**interval** *interval*] [**timeout** *timeout*]

Context <GLOBAL>
config>saa>test>type

Description Displays the hop-by-hop path for a destination MAC address within a VPLS.

The MAC traceroute operation is modeled after the IP traceroute utility which uses ICMP echo request and reply packets with increasing TTL values to determine the hop-by-hop route to a destination IP. The MAC traceroute command uses Alcatel-Lucent OAM packets with increasing TTL values to determine the hop-by-hop route to a destination MAC.

In a MAC traceroute, the originating device creates a MAC ping echo request packet for the MAC to be tested with increasing values of the TTL. The echo request packet is sent through the control plane or data plane and awaits a TTL exceeded response or the echo reply packet from the device with the destination MAC. The devices that reply to the echo request packets with the TTL exceeded and the echo reply are displayed.

When a **source** *ieee-address* value is specified and the source MAC address is locally registered within a split horizon group (SHG), then this SHG membership will be used as if the packet

originated from this SHG. In all other cases, SHG 0 (zero) will be used. Note that if the **mac-ping** is originated from a non-zero SHG, such packets will not go out to the same SHG.

If EMG is enabled, mac-trace will return only the first SAP in each chain.

Parameters **service** *service-id* **—** The Service ID of the service to diagnose or manage.

Values 1 — 2147483647

destination *ieee-address* **—** The destination MAC address to be traced.

size *octets* **—** The MAC OAM request packet size in octets, expressed as a decimal integer. The request payload is padded to the specified size with a 6 byte PAD header and a byte payload of 0xAA as necessary. If the octet size specified is less than the minimum packet, the minimum sized packet necessary to send the request is used.

Default No OAM packet padding.

Values 1 — 65535

min-ttl *vc-label-ttl* **—** The minimum TTL value in the VC label for the MAC trace test, expressed as a decimal integer.

Default 1

Values 1 — 255

max-ttl *vc-label-ttl* **—** The maximum TTL value in the VC label for the MAC trace test, expressed as a decimal integer.

Default 4

Values 1 — 255

send-control — Specifies the MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM request sent using the data plane.

return-control — Specifies the MAC OAM reply to a data plane MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM reply sent using the data plane.

source *ieee-address* **—** The source MAC address from which the OAM MAC request originates. By default, the system MAC address for the chassis is used.

Default The system MAC address.

Values Any unicast MAC value.

send-count *send-count* **—** The number of MAC OAM requests sent for a particular TTL value, expressed as a decimal integer.

Default 1

Values 1 — 100

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

sdp-ping

Syntax **sdp-ping** *orig-sdp-id* [**resp-sdp** *resp-sdp-id*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**timeout** *seconds*] [**interval** *seconds*] [**size** *octets*] [**count** *send-count*]

Context <GLOBAL>
config>saa>test>type

Description Tests SDPs for uni-directional or round trip connectivity and performs SDP MTU Path tests.

The **sdp-ping** command accepts an originating SDP-ID and an optional responding SDP-ID. The size, number of requests sent, message time-out and message send interval can be specified. All **sdp-ping** requests and replies are sent with PLP OAM-Label encapsulation, as a *service-id* is not specified.

For round trip connectivity testing, the **resp-sdp** keyword must be specified. If **resp-sdp** is not specified, a uni-directional SDP test is performed.

To terminate an **sdp-ping** in progress, use the CLI break sequence <Ctrl-C>.

An **sdp-ping** response message indicates the result of the **sdp-ping** message request. When multiple response messages apply to a single SDP Echo Request/Reply sequence, the response message with the highest precedence will be displayed. The following table displays the response messages sorted by precedence.

Result of Request	Displayed Response Message	Precedence
Request timeout without reply	Request Timeout	1
Request not sent due to non-existent *orig-sdp-id*	Orig-SDP Non-Existent	2
Request not sent due to administratively down *orig-sdp-id*	Orig-SDP Admin-Down	3
Request not sent due to operationally down *orig-sdp-id*	Orig-SDP Oper-Down	4

Result of Request	Displayed Response Message	Precedence
Request terminated by user before reply or timeout	Request Terminated	5
Reply received, invalid *origination-id*	Far End: Originator-ID Invalid	6
Reply received, invalid *responder-id*	Far End: Responder-ID Error	7
Reply received, non-existent *resp-sdp-id*	Far End: Resp-SDP Non-Existent	8
Reply received, invalid *resp-sdp-id*	Far End: Resp-SDP Invalid	9
Reply received, *resp-sdp-id* down (admin or oper)	Far-end: Resp-SDP Down	10
Reply received, No Error	Success	11

Parameters *orig-sdp-id* — The SDP-ID to be used by **sdp-ping,** expressed as a decimal integer. The far-end address of the specified SDP-ID is the expected *responder-id* within each reply received. The specified SDP-ID defines the encapsulation of the SDP tunnel encapsulation used to reach the far end. This can be IP/GRE or MPLS. If *orig-sdp-id* is invalid or administratively down or unavailable for some reason, the SDP Echo Request message is not sent and an appropriate error message is displayed (once the **interval** timer expires, sdp-ping will attempt to send the next request if required).

Values 1 — 17407

resp-sdp *resp-sdp-id* **—** Optional parameter is used to specify the return SDP-ID to be used by the far-end 7750 SR for the message reply for round trip SDP connectivity testing. If *resp-sdp-id* does not exist on the far-end 7750 SR, terminates on another 7750 SR different than the originating 7750 SR, or another issue prevents the far-end 7750 SR from using *resp-sdp-id*, the SDP Echo Reply will be sent using generic IP/GRE OAM encapsulaton. The received forwarding class (as mapped on the ingress network interface for the far end) defines the forwarding class encapsulation for the reply message.

Default null. Use the non-SDP return path for message reply.

Values 1 — 17407

fc *fc-name* **—** The **fc** parameter is used to indicate the forwarding class of the SDP encapsulation. The actual forwarding class encoding is controlled by the network egress DSCP or LSP-EXP mappings.

The DSCP or LSP-EXP mappings on the receive network interface controls the mapping back to the internal forwarding class used by the far-end 7750 SR that receives the message request. The egress mappings of the egress network interface on the far-end 7750 SR controls the forwarding class markings on the return reply message.

The DSCP or LSP-EXP mappings on the receive network interface controls the mapping of the message reply back at the originating 7750 SR. This is displayed in the response message output upon receipt of the message reply.

Default be

Values be, l2, af, l1, h2, ef, h1, nc

profile {**in** | **out**} **—** The profile state of the SDP encapsulation.

Default out

timeout *seconds* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A 'request timeout' message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

interval *seconds* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

size *octets* **—** The **size** parameter in octets, expressed as a decimal integer. This parameter is used to override the default message size for the **sdp-ping** request. Changing the message size is a method of checking the ability of an SDP to support a **path-mtu**. The size of the message does not include the SDP encapsulation, VC-Label (if applied) or any DLC headers or trailers.

When the OAM message request is encapsulated in an IP/GRE SDP, the IP 'DF' (Do Not Fragment) bit is set. If any segment of the path between the sender and receiver cannot handle the message size, the message is discarded. MPLS LSPs are not expected to fragment the message either, as the message contained in the LSP is not an IP packet.

Default 40

Values 40 — 9198

count *send-count* **—** The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

Special Cases **Single Response Connectivity Tests —** A single response sdp-ping test provides detailed test results.

Upon request timeout, message response, request termination, or request error the following local and remote information will be displayed. Local and remote information will be dependent upon SDP-ID existence and reception of reply.

Field	Description	Values
Request Result	The result of the **sdp-ping** request message.	Sent - Request Timeout
		Sent - Request Terminated
		Sent - Reply Received
		Not Sent - Non-Existent Local SDP-ID
		Not Sent - Local SDP-ID Down
Originating SDP-ID	The originating SDP-ID specified by **orig-sdp**.	*orig-sdp-id*
Originating SDP-ID Administrative State	The local administrative state of the originating SDP-ID. If the SDP-ID has been shutdown, Admin-Down is displayed. If the originating SDP-ID is in the no shutdown state, Admin-Up is displayed. If the *orig-sdp-id* does not exist, Non-Existent is displayed.	Admin-Up
		Admin-Down
		Non-Existent
Originating SDP-ID Operating State	The local operational state of the originating SDP-ID. If *orig-sdp-id* does not exist, N/A will be displayed.	Oper-Up
		Oper-Down
		N/A
Originating SDP-ID Path MTU	The local **path-mtu** for *orig-sdp-id*. If *orig-sdp-id* does not exist locally, N/A is displayed.	*orig-path-mtu*
		N/A
Responding SDP-ID	The SDP-ID requested as the far-end path to respond to the **sdp-ping** request. If **resp-sdp** is not specified, the responding 7750 SR will not use an SDP-ID as the return path and N/A will be displayed.	*resp-sdp-id*
		N/A

Field	Description	Values
Responding SDP-ID Path Used	Displays whether the responding 7750 SR used the responding SDP-ID to respond to the **sdp-ping** request. If *resp-sdp-id* is a valid, operational SDP-ID, it must be used for the SDP Echo Reply message. If the far-end uses the responding SDP-ID as the return path, Yes will be displayed. If the far-end does not use the responding SDP-ID as the return path, No will be displayed. If **resp-sdp** is not specified, N/A will be displayed.	Yes No N/A
Responding SDP-ID Administrative State	The administrative state of the responding SDP-ID. When *resp-sdp-id* is administratively down, Admin-Down will be displayed. When *resp-sdp-id* is administratively up, Admin-Up will be displayed. When *resp-sdp-id* exists on the far-end 7750 SR but is not valid for the originating 7750 SR, Invalid is displayed. When *resp-sdp-id* does not exist on the far-end 7750 SR, Non-Existent is displayed. When **resp-sdp** is not specified, N/A is displayed.	Admin-Down Admin-Up Invalid Non-Existent N/A
Responding SDP-ID Operational State	The operational state of the far-end SDP-ID associated with the return path for *service-id*. When a return path is operationally down, Oper-Down is displayed. If the return SDP-ID is operationally up, Oper-Up is displayed. If the responding SDP-ID is non-existent, N/A is displayed.	Oper-Up Oper-Down N/A
Responding SDP-ID Path MTU	The remote **path-mtu** for *resp-sdp-id*. If *resp-sdp-id* does not exist remotely, N/A is displayed	*resp-path-mtu* N/A
Local Service IP Address	The local system IP address used to terminate remotely configured SDP-IDs (as the SDP-ID **far-end** address). If an IP address has not been configured to be the system IP address, N/A is displayed.	*system-ip-addr* N/A
Local Service IP Interface Name	The name of the local system IP interface. If the local system IP interface has not been created, N/A is displayed.	*system-interface-name* N/A

Field	Description	Values
Local Service IP Interface State	The state of the local system IP interface. If the local system IP interface has not been created, Non-Existent is displayed.	Up
		Down
		Non-Existent
Expected Far End Address	The expected IP address for the remote system IP interface. This must be the **far-end** address configured for the *orig-sdp-id*.	*orig-sdp-far-end-addr*
		dest-ip-addr
		N/A
Actual Far End Address	The returned remote IP address. If a response is not received, the displayed value is N/A. If the far-end service IP interface is down or non-existent, a message reply is not expected.	*resp-ip-addr*
		N/A
Responders Expected Far End Address	The expected source of the originators SDP-ID from the perspective of the remote 7750 SR terminating the SDP-ID. If the far-end cannot detect the expected source of the ingress SDP-ID, N/A is displayed.	*resp-rec-tunnel-far-end-addr*
		N/A
Round Trip Time	The round trip time between SDP Echo Request and the SDP Echo Reply. If the request is not sent, times out or is terminated, N/A is displayed.	*delta-request-reply*
		N/A

Single Response Round Trip Connectivity Test Sample Output

```
A:router1> sdp-ping 10 resp-sdp 22 fc ef
Request Result: Sent - Reply Received
RTT: 30ms

Err  SDP-ID Info            Local     Remote
__   SDP-ID:                10        22
__   Administrative State:  Up        Up
__   Operative State:       Up        Up
__   Path MTU               4470      4470
__   Response SDP Used:               Yes

Err  System IP Interface Info
Local Interface Name: "ESR-System-IP-Interface (Up to 32 chars)…"
__   Local IP Interface State:          Up
__   Local IP Address:                  10.10.10.11
__   IP Address Expected By Remote:     10.10.10.11
__   Expected Remote IP Address:        10.10.10.10
__   Actual Remote IP Address:          10.10.10.10

Err   FC Mapping Info      Local        Remote
__    Forwarding Class     Assured      Assured
__    Profile              In           In
```

Multiple Response Connectivity Tests — When the connectivity test count is greater than one (1), a single line is displayed per SDP Echo Request send attempt.

The request number is a sequential number starting with 1 and ending with the last request sent, incrementing by one (1) for each request. This should not be confused with the *message-id* contained in each request and reply message.

A response message indicates the result of the message request. Following the response message is the round trip time value. If any reply is received, the round trip time is displayed.

After the last reply has been received or response timed out, a total is displayed for all messages sent and all replies received. A maximum, minimum and average round trip time is also displayed. Error response and timed out requests do not apply towards the average round trip time.

Multiple Response Round Trip Connectivity Test Sample Output

```
A:router1> sdp-ping 6 resp-sdp 101size 1514 count 5
Request        Response       RTT
----------     ----------     -------
     1         Success        10ms
     2         Success        15ms
     3         Success        10ms
     4         Success        20ms
     5         Success        5ms
Sent:    5    Received:    5
Min: 5ms       Max: 20ms      Avg: 12ms
```

vccv-ping

Syntax **vccv-ping** *sdp-id:vc-id* [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr* **pw-id** *pw-id*][**reply-mode** {**ip-routed**|**control-channel**}] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**ttl** *vc-label-ttl*]

Context oam
config>saa>test

Description This command configures a Virtual Circuit Connectivity Verification (VCCV) ping test. A vccv-ping test checks connectivity of a VLL inband. It checks to verify that the destination (target) PE is the egress for the Layer 2 FEC. It provides for a cross-check between the dataplane and the control plane. It is inband which means that the vccv-ping message is sent using the same encapsulation and along the same path as user packets in that VLL. The vccv-ping test is the equivalent of the lsp-ping test for a VLL service. The vccv-ping reuses an lsp-ping message format and can be used to test a VLL configured over both an MPLS and a GRE SDP.

Note that VCCV ping can be initiated on TPE or SPE. If initiated on the SPE, the **reply-mode** parameter must be used with the ip-routed value The ping from the TPE can have either values or can be omitted, in which case the default value is used.

If a VCCV ping is initiated from TPE to neighboring a SPE (one segment only) it is sufficient to only use the *sdpid:vcid* parameter. However, if the ping is across two or more segments, at least the *sdpId:vcId*, **src-ip-address** *ip-addr*, **dst-ip-address** *ip-addr*, **ttl** *vc-label-ttl*and **pw-id** *pw-id* parameters are used where:

- The *src-ip-address* is system IP address of the router proceeding destination router.
- The *pwid* is actually the VC ID of the last pseudowire segment.
- The *vc-label-ttl* must have a value equal or higher than the number of pseudowire segments .

Parameters *sdp-id:vc-id* — The VC ID of the pseudowire being tested must be indicated with this parameter. The VC ID needs to exist on the local router and the far-end peer needs to indicate that it supports VCCV to allow the user to send vccv-ping message.

Values 1 — 17407:1 — 2147483647

src-ip-address *ip-addr* **—** Specifies the source IP address.

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

dst-ip-address *ip-addr* **—** Specifies the destination IP address.

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

pw-id *pw-id* **—** Specifies the pseudowire ID to be used for performing a **vccv-ping** operation. The pseudowire ID is a non-zero 32-bit connection ID required by the FEC 128, as defined in RFE 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures.*

reply-mode {**ip-routed** | **control-channel**} **—** The reply-mode parameter indicates to the far-end how to send the reply message. The option ipv4 indicates a reply mode out-of-band using UDP IPv4. The option control-channel indicates a reply mode in-band using vccv control channel.

Default control-channel

fc *fc-name* **—** The **fc** parameter is used to indicate the forwarding class of the MPLS echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

The LSP-EXP mappings on the receive network interface controls the mapping back to the internal forwarding class used by the far-end 7x50 SR that receives the message request. The egress mappings of the egress network interface on the far-end router controls the forwarding class markings on the return reply message. The LSP-EXP mappings on the receive network interface controls the mapping of the message reply back at the originating SR.

Default be

Values be, l2, af, l1, h2, ef, h1, nc

profile {**in** | **out**} **—** The profile state of the MPLS echo request encapsulation.

Default out

timeout *seconds* **—** The timeout parameter, in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A 'request timeout' message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

interval *seconds* **—** The interval parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the interval is set to 1 second, and the timeout value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

size *octets* **—** The VCCV ping echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeroes to the specified size.

Default 88

Values 88 — 9198

count *send-count* **—** The number of messages to send, expressed as a decimal integer. The count parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message interval value must be expired before the next message request is sent.

Default 1

Values 1 — 100

ttl *vc-label-ttl* **—** Specifies the time-to-live value for the vc-label of the echo request message. The outer label TTL is still set to the default of 255 regardless of this value.

Sample Output

Ping from TPE to TPE:

```
*A:ALA-dut-b_a# oam vccv-ping 1:1 src-ip-address 5.5.5.5 dst-ip-address 3.3.3.3 pw-id
1 ttl 3
VCCV-PING 1:1 88 bytes MPLS payload
Seq=1, reply from 3.3.3.3 via Control Channel
       udp-data-len=32 rtt=10ms rc=3 (EgressRtr)

---- VCCV PING 1:1 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 10.0ms, avg = 10.0ms, max = 10.0ms, stddev < 10ms
```

Ping from TPE to SPE:

```
*A:ALA-dut-b_a# oam vccv-ping 1:1
VCCV-PING 1:1 88 bytes MPLS payload
Seq=1, reply from 4.4.4.4 via Control Channel
       udp-data-len=32 rtt<10ms rc=8 (DSRtrMatchLabel)

---- VCCV PING 1:1 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min < 10ms, avg < 10ms, max < 10ms, stddev < 10ms


*A:ALA-dut-b_a# oam vccv-ping 1:1 src-ip-address 4.4.4.4 dst-ip-address 5.5.5.5 ttl 2
pw-id 200
VCCV-PING 1:1 88 bytes MPLS payload
Seq=1, reply from 5.5.5.5 via Control Channel
       udp-data-len=32 rtt<10ms rc=8 (DSRtrMatchLabel)


---- VCCV PING 1:1 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min < 10ms, avg < 10ms, max < 10ms, stddev < 10ms
```

Ping from SPE (on single or multi-segment):

```
*A:ALA-dut-b_a# oam vccv-ping 4:200 reply-mode ip-routed
VCCV-PING 4:200 88 bytes MPLS payload
Seq=1, reply from 5.5.5.5 via IP
       udp-data-len=32 rtt<10ms rc=8 (DSRtrMatchLabel)

---- VCCV PING 4:200 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min < 10ms, avg < 10ms, max < 10ms, stddev < 10ms


*A:ALA-dut-b_a# oam vccv-ping 4:200 reply-mode ip-routed src-ip-address 5.5.5.5 dst-
ip-address 3.3.3.3 ttl 2 pw-id 1
VCCV-PING 4:200 88 bytes MPLS payload
Seq=1, reply from 3.3.3.3 via IP
       udp-data-len=32 rtt<10ms rc=3 (EgressRtr)

---- VCCV PING 4:200 Statistics ----
```

```
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min < 10ms, avg < 10ms, max < 10ms, stddev < 10ms
```

OAM SAA Commands

saa

Syntax **saa** *test-name* [**owner** *test-owner*] {**start** | **stop**}

Context oam

Description Use this command to start or stop an SAA test.

test-name — Name of the SAA test. The test name must already be configured in the **config>saa>test** context.

owner *test-owner* **—** Specifies the owner of an SAA operation up to 32 characters in length.

Values If a *test-owner* value is not specified, tests created by the CLI have a default owner "TiMOS CLI".

start — This keyword starts the test. A test cannot be started if the same test is still running.

A test cannot be started if it is in a shut-down state. An error message and log event will be generated to indicate a failed attempt to start an SAA testrun.

stop — This keyword stops a test in progress. A test cannot be stopped if it is not in progress. A log message will be generated to indicate that an SAA testrun has been aborted.

LDP Treetrace Commands

ldp-treetrace

Syntax **ldp-treetrace** {**prefix** *ip-prefix/mask*} [**max-ttl** *ttl-value*] [**max-path** *max-paths*] [**timeout** *timeout*] [**retry-count** *retry-count*] [**fc** *fc-name* [**profile** *profile*]]

Context oam

Description This command enables the context to configure LDP treetrace parameters to perform Alcatel-Lucent OAM tree trace test operations manually.

Parameters **prefix** *ip-prefix/mask* **—** Specifies the address prefix and subnet mask of the destination node.

max-ttl *max-label-ttl* **—** The maximum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer.

Default 30

Values 1 — 255

max-paths *max-paths* **—** The maximum number of paths for a ldp-treetrace test, expressed as a decimal integer.

Default 128

Values 1 — 255

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 3

Values 1 — 60

fc *fc-name* **—** The **fc** and **profile** parameters are used to indicate the forwarding class of the MPLS echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

The LSP-EXP mappings on the receive network interface controls the mapping back to the internal forwarding class used by the far-end 7750 SR that receives the message request. The egress mappings of the egress network interface on the far-end 7750 SR controls the forwarding class markings on the return reply message.

The LSP-EXP mappings on the receive network interface controls the mapping of the message reply back at the originating 7750 SR.

Default be

Values be, l2, af, l1, h2, ef, h1, nc

profile *profile* **—** The profile state of the MPLS echo request encapsulation.

Default out

Values in, out

retry-count *retry-count* **—** Specifies the maximum number of consecutive MPLS echo requests, expressed as a decimal integer that do not receive a reply before the trace operation fails for a given TTL.

Default 5

Values 1 — 255

ldp-treetrace

Syntax [**no**] **ldp-treetrace**

Context config>test-oam

Description This command enables the context to configure LDP treetrace parameters to perform Alcatel-Lucent OAM tree trace test operations manually**.**

The **no** form of the command disables the LDP treetrace parameters.

fc

Syntax **fc** *fc-name* [**profile** {**in** | **out**}]
no fc

Context config>test-oam>ldp-treetrace

Description This command configures forwarding class name and profile parameters. These parameters indicate the forwarding class of the MPLS echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

The LSP-EXP mappings on the receive network interface controls the mapping back to the internal forwarding class used by the far-end 7750 SR that receives the message request. The egress mappings of the egress network interface on the far-end 7750 SR controls the forwarding class markings on the return reply message.

The LSP-EXP mappings on the receive network interface controls the mapping of the message reply back at the originating 7750 SR.

Default **be**

Parameters *fc-name* — Specifies the forwarding class of the MPLS echo request packets.

Values be, l2, af, l1, h2, ef, h1, nc

profile {**in** | **out**}] **—** Specifies the profile value to be used with the forwarding class specified in the *fc-name* parameter.

path-discovery

Syntax **path-discovery**

Context config>test-oam>ldp-treetrace

Description This command enables the context to configure path discovery parameters.

interval

Syntax **interval** *minutes*
no interval

Context config>test-oam>ldp-treetrace>path-discovery

Description This command configures the time to wait before repeating the LDP Tree auto discovery process.

Default **60**

Parameters *minutes* — Specifies the number of minutes to wait before repeating the LDP Tree auto discovery process.

Values 60 — 1440

max-path

Syntax **max-path** *max-paths*

Context config>test-oam>ldp-treetrace>path-discovery

Description This command configures specifies the maximum number of paths that can be discovered for a selected IP address FEC.

Default **128**

Parameters *max-paths* — Specifies the tree discovery maximum path.

Values 1 — 128

max-ttl

Syntax **max-ttl** *ttl-value*

Context config>test-oam>ldp-treetrace>path-discovery

Description This command configures the maximum label time-to-live value for an LSP trace request during the tree discovery.

Default **30**

Parameters *ttl-value* — Specifies the maximum label time-to-live value for an LSP trace request during the tree discovery.

Values 1 — 255

policy-statement

Syntax **policy-statement** *policy-name* [...(up to 5 max)]

Context config>test-oam>ldp-treetrace>path-discovery

Description This command specifies policies to filter LDP imported address FECs.

Default **no policy-statement**

Parameters *policy-name* — Specifies the route policy name to filter LDP imported address FECs. 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.

retry-count

Syntax **retry-count** *retry-count*

Context config>oam-test>ldp-treetrace>path-discovery
config>oam-test>ldp-treetrace>path-probing

Description This command configures the path probing maximum number of failures.

Default **3**

Parameters *retry-count* — Specifies the maximum number of consecutive timeouts allowed before failing a path probe (ping).

Values 1 — 255

timeout

Syntax **timeout** *timeout*
no timeout

Context config>test-oam>ldp-treetrace>path-discovery

Description This command is used to override the default timeout value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default **30**

Parameters *timeout* — Specifies the timeout parameter, in seconds, within a range of 1 to 60, expressed as a decimal integer.

path-probing

Syntax **path-probing**

Context config>test-oam>ldp-treetrace

Description This command enables the context to configure path probing paramters.

interval

Syntax **interval** *minutes*
no interval

Context config>test-oam>ldp-treetrace>path-probing

Description This command configures the number of minutes to wait before repeating probing (pinging) a discovered path.

Default **1**

Parameters *minutes* — Specifies the number of minutes to probe all active ECMP paths for each LSP

Values 1 — 60

retry-count

Syntax **retry-count** *retry-count*

Context config>oam-test>ldp-treetrace>path-discovery
config>oam-test>ldp-treetrace>path-probing

Description This command configures the path probing maximum number of failures.

Default **3**

Parameters *retry-count* — Specifies the maximum number of consecutive timeouts allowed before failing a path probe (ping).

Values 1 — 255

timeout

Syntax **timeout** *timeout*
no timeout

Context config>test-oam>ldp-treetrace>path-probing

Description This command is used to override the default timeout value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default **1**

Parameters *timeout* — Specifies the timeout parameter, in minutes, with a range of 1 to 3 minutes, expressed as a decimal integer.

Show Commands

saa

Syntax **saa** [*test-name*] [**owner** *test-owner*]

Context show>saa

Description Use this command to display information about the SAA test.

If no specific test is specified a summary of all configured tests is displayed.

If a specific test is specified then detailed test results for that test are displayed for the last three occurrences that this test has been executed, or since the last time the counters have been reset via a system reboot or clear command.

Parameters *test-name* — Enter the name of the SAA test for which the information needs to be displayed. The test name must already be configured in the **config>saa>test** context.

This is an optional parameter.

owner *test-owner* — Specifies the owner of an SAA operation up to 32 characters in length.

Values 32 characters maximum.

Default If a *test-owner* value is not specified, tests created by the CLI have a default owner "TiMOS CLI".

Output **SAA Output —** The following table provides SAA field descriptions

Label	Description
Test Name	The name of the test.
Owner Name	The owner of the test.
Administrative status	Enabled or disabled.
Test type	The type of test configured.
Test runs since last clear	The total number of tests performed since the last time the tests were cleared.
Number of failed tests run	The total number of tests that failed.
Last test run	The last time a test was run.

Sample Output

The following displays an SAA test result:

```
*A:SR-3>config>saa>test$ show saa
```

```
===============================================================================
SAA Test Information
===============================================================================
Test name                    : test1
Owner name                   : TiMOS CLI
Administrative status        : Disabled
Test type                    : cpe-ping service 1 source 192.168.1.1
                               destination 192.168.1.1
Test runs since last clear   : 0
Number of failed test runs   : 0
Last test result             : Undetermined
-------------------------------------------------------------------------------
Threshold
Type        Direction Threshold  Value      Last Event          Run #
-------------------------------------------------------------------------------
Jitter-in   Rising    None       None       Never               None
            Falling   None       None       Never               None
Jitter-out  Rising    None       None       Never               None
            Falling   None       None       Never               None
Jitter-rt   Rising    None       None       Never               None
            Falling   None       None       Never               None
Latency-in  Rising    None       None       Never               None
            Falling   None       None       Never               None
Latency-out Rising    None       None       Never               None
            Falling   None       None       Never               None
Latency-rt  Rising    None       None       Never               None
            Falling   None       None       Never               None
Loss-in     Rising    None       None       Never               None
            Falling   None       None       Never               None
Loss-out    Rising    None       None       Never               None
            Falling   None       None       Never               None
Loss-rt     Rising    None       None       Never               None
            Falling   None       None       Never               None

===============================================================================
*A:SR-3>config>saa>test$
```

ldp-treetrace

Syntax **ldp-treetrace** [**prefix** *ip-prefix/mask*] [**detail**]

Context show>test-oam

Description This command displays OAM LDP treetrace information.

Parameters **prefix** *ip-prefix/mask* **—** Specifies the address prefix and subnet mask of the destination node.

detail — Displays detailed information.

Sample Output

```
*A:ALA-48# show test-oam ldp-treetrace
Admin State              : Up             Discovery State      : Done
Discovery-intvl (min)    : 60             Probe-intvl (min)    : 2
Probe-timeout (min)      : 1              Probe-retry          : 3
Trace-timeout (sec)      : 60             Trace-retry          : 3
```

```
Max-TTL                  : 30             Max-path             : 128
Forwarding-class (fc)    : be             Profile              : Out
Total Fecs               : 400            Discovered Fecs      : 400
Last Discovery Start     : 12/19/2006 05:10:14
Last Discovery End       : 12/19/2006 05:12:02
Last Discovery Duration  : 00h01m48s
Policy1                  : policy-1
Policy2                  : policy-2

*A:ALA-48# show test-oam ldp-treetrace detail
Admin State              : Up             Discovery State      : Done
Discovery-intvl (min)    : 60             Probe-intvl (min)    : 2
Probe-timeout (min)      : 1              Probe-retry          : 3
Trace-timeout (sec)      : 60             Trace-retry          : 3
Max-TTL                  : 30             Max-path             : 128
Forwarding-class (fc)    : be             Profile              : Out
Total Fecs               : 400            Discovered Fecs      : 400
Last Discovery Start     : 12/19/2006 05:10:14
Last Discovery End       : 12/19/2006 05:12:02
Last Discovery Duration  : 00h01m48s
Policy1                  : policy-1
Policy2                  : policy-2
===============================================================================
Prefix (FEC) Info
===============================================================================
Prefix             Path Last                  Probe  Discov    Discov
                   Num  Discovered            State  State     Status
-------------------------------------------------------------------------------
11.11.11.1/32      54   12/19/2006 05:10:15   OK     Done      OK
11.11.11.2/32      54   12/19/2006 05:10:15   OK     Done      OK
11.11.11.3/32      54   12/19/2006 05:10:15   OK     Done      OK
……
 14.14.14.95/32     72   12/19/2006 05:11:13   OK     Done      OK
14.14.14.96/32     72   12/19/2006 05:11:13   OK     Done      OK
14.14.14.97/32     72   12/19/2006 05:11:15   OK     Done      OK
14.14.14.98/32     72   12/19/2006 05:11:15   OK     Done      OK
14.14.14.99/32     72   12/19/2006 05:11:18   OK     Done      OK
14.14.14.100/32    72   12/19/2006 05:11:20   OK     Done      OK
===============================================================================

Legend: uP - unexplored paths, tO - trace request timed out
        mH - max hop exceeded, mP - max path exceeded
        nR - no internal resource

*A:ALA-48# show test-oam ldp-treetrace prefix 12.12.12.10/32
Discovery State  : Done                Last Discovered  : 12/19/2006 05:11:02
Discovery Status : ' OK '
Discovered Paths : 54                  Failed Hops      : 0
Probe State      : OK                  Failed Probes    : 0


*A:ALA-48# show test-oam ldp-treetrace prefix 12.12.12.10/32 detail
*A:ALA-48# show test-oam ldp-treetrace prefix 12.12.12.10/32  detail
Discovery State  : Done                Last Discovered  : 12/19/2006 05:11:02
Discovery Status : ' OK '
Discovered Paths : 54                  Failed Hops      : 0
Probe State      : OK                  Failed Probes    : 0
===============================================================================
Discovered Paths
===============================================================================
PathDest           Egr-NextHop         Remote-RtrAddr      Discovery-time
```

```
  DiscoveryTtl        ProbeState          ProbeTmOutCnt       RtnCode
-------------------------------------------------------------------------------
127.1.0.5           10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  0                   EgressRtr
127.1.0.9           10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  0                   EgressRtr
127.1.0.15          10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  0                   EgressRtr
127.1.0.19          10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  0                   EgressRtr
127.1.0.24          10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  0                   EgressRtr
127.1.0.28          10.10.1.2           12.12.12.10         12/19/2006 05:11:01

……….

127.1.0.252         10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  0                   EgressRtr
127.1.0.255         10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  0                   EgressRtr
===============================================================================
*A:ALA-48#
```

Clear Commands

saa

Syntax **saa-test** [*test-name*] [**owner** *test-owner*]

Context clear

Description Clear the SAA results for the latest and the history for this test. If the test name is omitted, all the results for all tests are cleared.

Parameters *test-name* — Name of the SAA test. The test name must already be configured in the **config>saa>test** context.

owner *test-owner* **—** Specifies the owner of an SAA operation up to 32 characters in length.

Default If a *test-owner* value is not specified, tests created by the CLI have a default owner "TiMOS CLI".

Debug Commands

lsp-ping-trace

Syntax **lsp-ping-trace** [**tx** | **rx** | **both**] [**raw** | **detail**]
no lsp-ping-trace

Context **debug>oam**

Description This command enables debugging for lsp-ping.

Parameters **tx**|**rx**|**both —** Specifies to enable LSP ping debugging for TX, RX, or both RX and TX for the for debug direction.

raw | detail — Displays output for the for debug mode.

Tools Configuration Commands

Generic Commands

tools

Syntax **tools**

Context <root>

Description The context to enable useful tools for debugging purposes.

Default **none**

Parameters **dump —** Enables dump tools for the various protocols.

perform — Enables tools to perform specific tasks.

Dump Commands

dump

Syntax **dump** *router-name*

Context tools

Description The context to display information for debugging purposes.

Default **none**

Parameters *router-name* — Specify a router name, up to 32 characters in length.

Default Base

lag

Syntax **lag lag-id** *lag-id*

Context tools>dump

Description This tool displays LAG information.

Parameters *lag-id* — Specify an existing LAG id.

Values 1 — 200

```
ALA-12>tools>dump# lag lag-id 1
Port state      : Ghost
Selected subgrp : 1
NumActivePorts  : 0
ThresholdRising : 0
ThresholdFalling: 0
IOM bitmask     : 0
Config MTU      : 1514
Oper. MTU       : 1514
Bandwidth       : 100000
ALA-12>tools>dump#
```

ldp-treetrace

Syntax **ldp-treetrace** {**prefix** *ip-prefix/mask*| **manual-prefix** *ip-prefix/mask*}[**path-destination** *ip-address*] [**trace-tree**]

Context tools>dump

Description This command displays TreeTrace information.

Parameters **prefix** *ip-prefix/mask* — Specifies the IP prefix and host bits.

Values

host bits:	must be 0
mask:	0 — 32

persistence

Syntax persistence

Context tools>dump

Description This command enables the context to display persistence information for debugging purposes.

submgt

Syntax **submgt** [**record** *record-key*]

Context tools>dump>persistence

Description This command displays subscriber management persistence info.

summary

Syntax **summary**

Context tools>dump>persistence

Description The context to display persistence summary information for debugging purposes.

Sample Output

```
A:ALA-B# tools dump persistence summary
===============================================================================
Persistence Summary on Slot A
===============================================================================
Client             Location            Entries in use    Status
-------------------------------------------------------------------------------
xxxxxx            cf1:\l2_dhcp.pst    200               ACTIVE
-------------------------------------------------------------------------------

===============================================================================
Persistence Summary on Slot B
===============================================================================
Client             Location            Entries in use    Status
-------------------------------------------------------------------------------
xxxxxx              cf1:\l2_dhcp.pst    200               ACTIVE
-------------------------------------------------------------------------------
A:ALA-B#
```

ppp

Syntax **ppp** *port-id*

Context tools>dump

Description This command displays PPP information for a port.

Default **none**

Parameters *port-id* — Specify the port ID.

Syntax: *port-id* *slot/mda/port*[.*channel*]
aps-id aps-*group-id*[.*channel*]
aps keyword
group-id 1 — 64
bundle bundle-*type-slot/mda*.*bundle-num*
bundle keyword
type ima, ppp
bundle-num 1 — 128

system-resources

Syntax **system-resources** *slot-number*

Context tools>dump

Description This command displays system resource information.

Default **none**

slot-number — Specify a specific slot to view system resources information.

Service Commands

service

Syntax **service**

Context tools>dump

Description Use this command to configure tools to display service dump information.

base-stats

Syntax **base-stats** [**clear**]

Context tools>dump>service

Description Use this command to display internal service statistics.

Default **none**

Parameters **clear —** Clears stats after reading.

iom-stats

Syntax **iom-stats** [**clear**]

Context tools>dump>service

Description Use this command to display IOM message statistics.

Default **none**

Parameters **clear —** Clears stats after reading.

l2pt-diags

Syntax **l2pt-diags**
l2pt-diags clear
l2pt-diags detail

Context tools>dump>service

Description Use this command to display L2pt diagnostics.

Default **none**

Parameters **clear —** Clears the diags after reading.

detail — Displays detailed information.

Sample Output

```
A:ALA-48>tools>dump>service# l2pt-diags
[ l2pt/bpdu error diagnostics ]
 Error Name       | Occurence   | Event log
 -----------------+-------------+-----------------------------
[ l2pt/bpdu forwarding diagnostics ]

 Rx Frames   | Tx Frames   | Frame Type
 ------------+-------------+---------------------------------
A:ALA-48>tools>dump>service#


A:ALA-48>tools>dump>service# l2pt-diags detail
[ l2pt/bpdu error diagnostics ]
 Error Name       | Occurence   | Event log
 -----------------+-------------+-----------------------------
[ l2pt/bpdu forwarding diagnostics ]

 Rx Frames   | Tx Frames   | Frame Type
 ------------+-------------+---------------------------------
[ l2pt/bpdu config diagnostics ]

 WARNING - service 700 has l2pt termination enabled on all access points :
           consider translating further down the chain or turning it off.
 WARNING - service 800 has l2pt termination enabled on all access points :
           consider translating further down the chain or turning it off.
 WARNING - service 9000 has l2pt termination enabled on all access points :
           consider translating further down the chain or turning it off.
 WARNING - service 32806 has l2pt termination enabled on all access points :
           consider translating further down the chain or turning it off.
 WARNING - service 90001 has l2pt termination enabled on all access points :
           consider translating further down the chain or turning it off.
A:ALA-48>tools>dump>service#
```

radius-discovery

Syntax **radius-discovery** [**svc-id** *service-id*]

Context tools>dump>service

Description Use thie command to display RADIUS Discovery membership information.

Sample Output

```
A:ALA-48# tools dump service radius-discovery
-------------------------------------------------------------
Service Id 103  Vpn Id 103  UserName 901:103 (Vpn-Id)  PolicyName RAD_Disc for Ser-
vice 103
Waiting for Session Timeout (Polling 60), Seconds in State 17
-------------------------------------------------------------------------------
      SdpId      Vcid  Deliver     Ip Addr        VcType       Mode      Split Horizon
-------------------------------------------------------------------------------
          3       103    LDP    10. 20.  1.  3     Ether      Spoke
          4       103    LDP    10. 20.  1.  2     Ether      Spoke
 -------------------------------------------------------------
```

```
A:ALA-48#
```

vpls-fdb-stats

Syntax **vpls-fdb** [**clear**]

Context tools>dump>service

Description Use this command to display VPLS FDB statistics.

Default **none**

Parameters **clear —** Clears stats after reading.

vpls-mfib-stats

Syntax **vpls-mfib-stats** [**clear**]

Context tools>dump>service

Description Use this command to display VPLS MFIB statistics.

Default **none**

Parameters **clear —** Clears stats after reading.

Router Commands

router

Syntax **router** *router-instance*

Context tools>dump
tools>perform

Description This command enables tools for the router instance.

Default **none**

Parameters **router** *router-instance* — Specifies the router name or service ID.

Values *router-name*: Base , management
service-id: 1 — 2147483647

Default Base

dhcp

Syntax **dhcp**

Context tools>dump>router

Description This command enables the context to configure dump router tools for DHCP.

group-if-mapping

Syntax **group-if-mapping** [**clear**]

Context tools>dump>router>dhcp

Description This command dumps group interface mapping information stored in by the DHCP cache for the Routed CO model of operation.

group-if-stats

Syntax **group-if-stats** [**clear**]

Context tools>dump>router>dhcp

Description This command dumps group interface statistics information about the DHCP cache for the Routed CO model of operation.

lag

Syntax **lag**

Context tools>perform

Description This command configures tools to control LAG.

clear-force

Syntax **clear-force all-mc**
clear-force lag-id *lag-id* [**sub-group** *sub-group-id*]
clear-force peer-mc *ip-address*

Context tools>perform>lag

Description This command clears a forced status.

Parameters **all-mc —**

lag-id *lag-id* — Specify an existing LAG id.

Values 1 — 200

sub-group *sub-group-id* —

peer-mc *ip-address* —

force

Syntax **force all-mc** {**active** | **standby**}
force lag-id *lag-id* [**sub-group** *sub-group-id*] {**active** | **standby**}
force peer-mc *peer-ip-address* {**active** | **standby**}

Context tools>perform>lag

Description This command forces an active or standy status.

Parameters **active —** If **active** is selected, then all drives on the active CPM are forced.

standby — If **standby** is selected, then all drives on the standby CPM are forced.

all-mc —

lag-id *lag-id* **—** Specify an existing LAG id.

Values 1 — 200

sub-group *sub-group-id* **—**

peer-mc *peer-ip-address* **—**

log

Syntax **log**

Context tools>perform

Description Tools for event logging.

test-event

Syntax **test-event**

Context tools>perform>log

Description Generates a test event.

ldp

Syntax **ldp**

Context tools>dump>router

Description This command enables dump tools for LDP.

Default **none**

interface

Syntax **interface** [*ip-int-name* | *ip-address*]

Context tools>dump>router>ldp

Description This command displays information for an LDP interface.

Default **none**

Parameters *ip-int-name* — Specifies the interface name.

ip-address — Specifies the IP address.

peer

Syntax **peer** *ip-address*

Context tools>dump>router>ldp

Description This command displays information for an LDP peer.

Default **none**

Parameters *ip-address* — Specifies the IP address.

fec

Syntax **fec prefix** [ip-prefix/mask]
fec vc-type {**ethernet**|**vlan**} **vc-id** *vc-id*

Context tools>dump>router>ldp

Description This command displays information for an LDP FEC.

Default **none**

Parameters *ip-prefix/mask* — Specifies the IP prefix and host bits.

Values

host bits:	must be 0
mask:	0 — 32

vc-type — Specifies the VC type signaled for the spoke or mesh binding to the far end of an SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the binding's VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

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

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

vc-id — Specifies the virtual circuit identifier.

Values 1 — 4294967295

instance

Syntax **instance**

Context tools>dump>router>ldp

Description This command displays information for an LDP instance.

memory-usage

Syntax **memory-usage**

Context tools>dump>router>ldp

Description This command displays memory usage information for LDP.

Default **none**

session

Syntax **session** [*ip-address* |*:label space*] [*connection* | *peer* | *adjacency*]

Context tools>dump>router>ldp

Description This command displays information for an LDP session.

Default **none**

Parameters *ip-address* — Specifies the IP address of the LDP peer.

label-space — Specifies the label space identifier that the router is advertising on the interface.

connection — Displays connection information.

peer — Displays peer information.

adjacency — Displays hello adjacency information.

sockets

Syntax **sockets**

Context tools>dump>router>ldp

Description This command displays information for all sockets being used by the LDP protocol.

Default **none**

timers

Syntax **timers**

Context tools>dump>router>ldp

Description This command displays timer information for LDP.

Default **none**

mpls

Syntax **mpls**

Context tools>dump>router

Description This command enables the context to display MPLS information.

Default **none**

ftn

Syntax **ftn**

Context tools>dump>router>mpls

Description This command displays FEC-to-NHLFE (FTN) dump information for MPLS. (NHLFE is the acronym for Next Hop Label Forwarding Entry.)

Default **none**

ilm

Syntax **ilm**

Context tools>dump>router>mpls

Description This command displays incoming label map (ILM) information for MPLS.

Default none

lspinfo

Syntax **lspinfo**

Context tools>dump>router>mpls

Description This command displays label-switched path (LSP) information for MPLS.

Default none

memory-usage

Syntax **memory-usage**

Context tools>dump>router>mpls

Description This command displays memory usage information for MPLS.

Default none

ospf

Syntax **ospf [ospf-instance]**

Context tools>dump>router

Description This command enables the context to display tools information for OSPF.

Parameters **ospf-instance —** OSPF instance.

Values 1 - 4294967295

Default none

ospf3

Syntax **ospf3**

Context tools>dump>router

Description This command enables the context to display tools information for OSPF3.

Default none

abr

Syntax **abr** [**detail**]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays area border router (ABR) information for OSPF.

Default **none**

Parameters **detail —** Displays detailed information about the ABR.

asbr

Syntax **asbr** [**detail**]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays autonoumous system border router (ASBR) information for OSPF.

Default **none**

Parameters **detail —** Displays detailed information about the ASBR.

bad-packet

Syntax **bad-packet** [*interface-name*]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays information about bad packets for OSPF.

Default **none**

Parameters *interface-name* — Display only the bad packets identified by this interface name.

leaked-routes

Syntax **leaked-routes** [**summary** | **detail**}

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays information about leaked routes for OSPF.

Default **summary**

Parameters **summary —** Display a summary of information about leaked routes for OSPF.

detail — Display detailed information about leaked routes for OSPF.

memory-usage

Syntax **memory-usage** [**detail**]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays memory usage information for OSPF.

Default **none**

Parameters **detail —** Displays detailed information about memory usage for OSPF.

request-list

Syntax **request-list** [**neighbor** *ip-address*] [**detail**]
request-list virtual-neighbor *ip-address* **area-id** *area-id* [**detail**]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays request list information for OSPF.

Default **none**

Parameters **neighbor** *ip-address* — Display neighbor information only for neighbor identified by the IP address.

detail — Displays detailed information about the neighbor.

virtual-neighbor *ip-address* — Displays information about the virtual neighbor identified by the IP address.

area-id *area-id* — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

retransmission-list

Syntax **retransmission-list** [**neighbor** *ip-address*] [**detail**]
retransmission-list virtual-neighbor *ip-address* **area-id** *area-id* [**detail**]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays dump retransmission list information for OSPF.

Default **none**

Parameters **neighbor** *ip-address* — Display neighbor information only for neighbor identified by the IP address.

detail — Displays detailed information about the neighbor.

virtual-neighbor *ip-address* — Displays information about the virtual neighbor identified by the IP address.

area-id *area-id* — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

route-summary

Syntax **route-summary**

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays dump route summary information for OSPF.

Default **none**

route-table

Syntax **route-table** [**type**] [**detail**]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays dump information about routes learned through OSPF.

Default **none**

Parameters **type —** Specify the type of route table to display information.

Values intra-area, inter-area, external-1, external-2, nssa-1, nssa-2

detail — Displays detailed information about learned routes.

te-database

Syntax **te-database** [**router** | **link**] [**summary** | **detail**]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays traffic engineering database information for OSPF.

Default **none**

Parameters **router —** Display te-database information for the router.

link — Display te-database information for the link.

summary — Displays a summary of information about the specified te-database.

detail — Displays detailed information about the specified te-database.

pim

Syntax **pim**

Context tools>dump>router

Description This command enables the context to display PIM information.

iom-failures

Syntax **iom-failures** [**detail**]

Context tools>dump>router>pim

Description This command displays information about failures in programming IOMs.

Parameters *detail* — Displays detailed information about IOM failures.

rsvp

Syntax rsvp

Context tools>dump>router

Description This command enables the context to display RSVP information.

Default **none**

psb

Syntax **psb** [**endpoint** *endpoint-address*] [**sender** *sender-address*] [**tunnelid** *tunnel-id*] [**lspid** *lsp-id*]

Context tools>dump>router>rsvp

Description This command displays path state block (PSB) information for RSVP.

When a PATH message arrives at an LSR, the LSR stores the label request in the local PSB for the LSP. If a label range is specified, the label allocation process must assign a label from that range.

The PSB contains the IP address of the previous hop, the session, the sender, and the TSPEC. This information is used to route the corresponding RESV message back to LSR 1.

Default **none**

Parameters **endpoint** *endpoint-address* — Specifies the IP address of the last hop.

sender *sender-address* — Specifies the IP address of the sender.

tunnelid *tunnel-id* — Specifies the SDP ID.

Values 0 — 4294967295

lspid *lsp-id* — Specifies the label switched path that is signaled for this entry.

Values 1 — 65535

rsb

Syntax **rsb** [**endpoint** *endpoint-address*] [**sender** *sender-address*] [**tunnelid** *tunnel-id*] [**lspid** *lsp-id*]

Context tools>dump>router>rsvp

Description This command displays RSVP Reservation State Block (RSB) information.

Default **none**

Parameters **endpoint** *endpoint-address* — Specifies the IP address of the last hop.

sender *sender-address* — Specifies the IP address of the sender.

tunnelid *tunnel-id* — Specifies the SDP ID.

Values 0 — 4294967295

lspid *lsp-id* — Specifies the label switched path that is signaled for this entry.

Values 1 — 65535

tcsb

Syntax **tcsb** [**endpoint** *endpoint-address*] [**sender** *sender-address*] [**tunnelid** *tunnel-id*] [**lspid** *lsp-id*]

Context tools>dump>router>rsvp

Description This command displays RSVP traffic control state block (TCSB) information.

Default **none**

Parameters **endpoint** *endpoint-address* — xxx

sender *sender-address* — xxx

tunnelid *tunnel-id* — xxx

Values 0 — 4294967295

lspid *lsp-id* — Specifies the label switched path that is signaled for this entry.

Values 1 — 65535

web-rd

Syntax **web-rd**

Context tools>dump>router

Description This command enables the context to display tools for web redirection.

http-client

Syntax **http-client** [*ip-prefix/mask*]

Context tools>dump>router>web-rd

Description This command displays the HTTP client hash table.

Parameters *ip-prefix/mask* — Specifies the IP prefix and host bits.

Values		
	host bits:	must be 0
	mask:	0 — 32

Performance Tools

perform

Syntax perform

Context tools

Description This command enables the context to enable tools to perform specific tasks.

Default **none**

cron

Syntax **cron**

Context tools>perform

Description This command enables the context to perform CRON (scheduling) control operations.

Default **none**

action

Syntax **action**

Context tools>perform>cron

Description This command enables the context to stop the execution of a script started by CRON action. See the **stop** command.

stop

Syntax **stop** [*action-name*] [**owner** *action-owner*] [**all**]

Context tools>perform>cron>action

Description This command stops execution of a script started by CRON action.

Parameters *action-name* — Specifies the action name.

Values Maximum 32 characters.

owner *action-owner* **—** Specifies the owner name.

Default TiMOS CLI

all **—** Specifies to stop all CRON scripts.

tod

Syntax **tod**

Context tools>perform>cron

Description This command enables the context for tools for controlling time-of-day actions.

Default **none**

re-evaluate

Syntax **re-evaluate**

Context tools>perform>cron

Description This command enables the context to re-evaluate the time-of-day state.

Default **none**

customer

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

Context tools>perform>cron>tod>re-eval

Description This command re-evaluates the time-of-day state of a multi-service site.

Parameters *customer-id* — Specify an existing customer ID.

Values 1 — 2147483647

site *customer-site-name* **—** Specify an existing customer site name.

filter

Syntax **filter** *filter-type* [*filter-id*]

Context tools>perform>cron>tod>re-eval

Description This command re-evaluates the time-of-day state of a filter entry.

Parameters *filter-type* — Specify the filter type.

Values ip-filter, ipv6-filter, mac-filter

filter-id — Specify an existing filter ID.

Values 1 — 65535

service

Syntax **service id** *service-id* [**sap** *sap-id*]

Context tools>perform>cron>tod>re-eval

Description This command re-evaluates the time-of-day state of a SAP.

Parameters **id** *service-id* — Specify the an existing service ID.

Values 1 — 2147483647

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

Values

sap-id:	null	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]
	dot1q	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]:*qtag1*
	qinq	[*port-id* | *bundle-id* | *lag-id*]:*qtag1.qtag2*
	atm	[*port-id* | *bundle-id*][:*vpi/vci* | *vpi* | *vpi1.vpi2*]
	frame	[*port-id* | *bundle-id*]:*dlci*
	cisco-hdlc	*slot/mda/port.channel*
	ima-grp	*bundle-id*[:*vpi/vci* | *vpi* | *vpi1.vpi2*]
	port-id	*slot/mda/port*[.*channel*]
	aps-id	aps-*group-id*[.*channel*]
		aps: keyword
		group-id: 1 — 64
	bundle-*type-slot/mda.bundle-num*	
		bundle: keyword
		type: ima, ppp
		bundle-num: 1 — 128
	ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*
		ccag: keyword
		id: 1 — 8
		path-id: a, b
		cc-type: .sap-net, .net-sap
		cc-id: 0 — 4094
	lag-id	lag-*id*
		lag: keyword
		id: 1 — 64
	qtag1	0 — 4094
	qtag2	*, 0 — 4094
	vpi	NNI 0 — 4095
		UNI 0 — 255]
	vci	1, 2, 5 — 65535
	dlci	16 — 1022

tod-suite

Syntax **tod-suite** *tod-suite-name*

Context tools>perform>cron>tod>re-eval

Description This command re-evaluates the time-of-day state for the objects referring to a tod-suite.

Parameters *tod-suite-name* — Specify an existing TOD name.

aps

Syntax **aps**

Context tools>perform

Description This command enables the context to perform Automated Protection Switching (APS) operations.

clear

Syntax **clear** *aps-id* {**protect** | **working**}

Context tools>perform>aps

Description This command removes all Automated Protection Switching (APS) operational commands.

Parameters *aps-id* — This option clears a specific APS on un-bundled SONET/SDH ports.

protect — This command clears a physical port that is acting as the protection circuit for the APS group.

working — This command clears a physical port that is acting as the working circuit for this APS group.

exercise

Syntax **exercise** *aps-id* {**protect** | **working**}

Context tools>perform>aps

Description This command performs an exercise request on the protection or working circuit.

Parameters *aps-id* — This option clears a specific APS on un-bundled SONET/SDH ports.

protect — This command performs an exercise request on the port that is acting as the protection circuit for the APS group.

working — This command performs an exercise request on the port that is acting as the working circuit for this APS group.

force

Syntax **force** *aps-id* {**protect** | **working**}

Context tools>perform>aps

Description This command forces a switch to either the protect or working circuit

Parameters *aps-id* — This option clears a specific APS on un-bundled SONET/SDH ports.

protect — This command clears a physical port that is acting as the protection circuit for the APS group.

working — This command clears a physical port that is acting as the working circuit for this APS group.

lockout

Syntax **lockout** *aps-id*

Context tools>perform>aps

Description This command locks out the protection circuit.

Parameters *aps-id* — Automated Protection Switching ID

Values 1 — 64

request

Syntax **request** *aps-id* {**protect** | **working**}

Context tools>perform>aps

Description This command requests a manual switch to protection or working circuit.

Parameters *aps-id* — This option clears a specific APS on un-bundled SONET/SDH ports.

protect — This command requests a manual switch to a port that is acting as the protection circuit for the APS group.

working — This command requests a manual switch to a port that is acting as the working circuit for this APS group.

consistency

Syntax **consistency**

Context tools>perform>router

Description This command performs route table manager (RTM) consistency checks.

Default **none**

isis

Syntax **isis**

Context tools>perform>router

Description This command enables the context to configure tools to perform certain ISIS tasks.

run-manual-spf

Syntax **run-manual-spf**

Context tools>perform>router>isis

Description This command runs the Shortest Path First (SPF) algorithm.

mpls

Syntax **mpls**

Context tools>perform>router

Description This command enables the context to perform specific MPLS tasks.

Default **none**

cspf

Syntax **cspf to** *ip-addr* [**from** *ip-addr*] [**bandwidth** *bandwidth*] [**include-bitmap** *bitmap*] [**exclude-bitmap** *bitmap*] [**hop-limit** *limit*] [**exclude-address** *ip-addr* [*ip-addr* ...(up to 8 max)]]

Context tools>perform>router>mpls

Description This command computes a CSPF path with specified user constraints.

Default **none**

Parameters **to** *ip-addr* — Specify the destination IP address.

from *ip-addr* — Specify the originating IP address.

bandwidth *bandwidth* — xxx

include-bitmap *bitmap* — xxx

exclude-bitmap *bitmap* — xxx

hop-limit *limit* — xxx

exclude-address *ip-addr* — xxx

resignal

Syntax **resignal lsp** *lsp-name* **path** *path-name*

Context tools>perform>router>mpls

Description Use this command to resignal a specific LSP path.

Default **none**

Parameters **lsp** *lsp-name* **—** Specifies the name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique.

path *path-name* **—** Specifies the name for the LSP path up, to 32 characters in length.

trap-suppress

Syntax **trap-suppress** [*number-of-traps*] [*time-interval*]

Context tools>perform>router>mpls

Description This command modifies thresholds for trap suppression.

Default **none**

Parameters *number-of-traps* — Specify the number of traps in multiples of 100. An error messages is generated if an invalid value is entered.

Values 100 to 1000

time-interval — Specify the timer interval in seconds.

Values 1 to 300

ospf

Syntax ospf

Context tools>perform>router

Description This command enables the context to perform specific OSPF tasks.

Default **none**

ospf3

Syntax ospf3

Context tools>perform>router

Description This command enables the context to perform specific OSPF3 tasks.

Default **none**

refresh-lsas

Syntax **refresh-lsas** [*lsa-type*] [*area-id*]

Context tools>perform>router>ospf

Description This command refreshes LSAs for OSPF.

Default **none**

Parameters *lsa-type* — Specify the LSA type using allow keywords.

Values Keywords: router, network, summary, asbr, extern, nssa, opaque

area-id — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer

Values 0 to 4294967295.

run-manual-spf

Syntax **run-manual-spf** *externals-only*

Context tools>perform>router>ospf

Description This command runs the Shortest Path First (SPF) algorithm.

Default **none**

Parameters **externals-only —** Specify the route preference for OSPF external routes.

security

Syntax **security**

Context tools>perform

Description This command provides tools for testing security.

authentication-server-check

Syntax **authentication-server-check** *server-address ip-address* [**port** *port*] **user-name** *DHCP client user name* **password** *password* **secret** *key* [**source-address** *ip-address*] [**timeout** *seconds*] [**router** *router-instance*]

Context tools>perform>security

Description This command checks connection to the RADIUS server.

Parameters **router** *router-instance* **—** Specifies the router name or service ID.

Values *router-name*: Base , management
service-id: 1 — 2147483647

Default Base

service

Syntax **services**

Context tools>perform

Description This command enables the context to configure tools for services.

egress-multicast-group

Syntax **egress-multicast-group** *group-name*

Context tools>perform>service

Description This command enables the context to configure tools for egress multicast groups.

Parameters *group-name* — Specify an existing group name.

force-optimize

Syntax **force-optimize**

Context tools>perform>service>egress-multicast-group

Description This command optimizes the chain length.

id

Syntax **id** *service-id*

Context tools>perform>service

Description This command enables the context to configure tools for a specific service.

Parameters *service-id* — Specify an existing service ID.

Values 1 — 2147483647

endpoint

Syntax **endpoint** *endpoint-name*

Context tools>perform>service>id

Description This command enables the context to configure tools for a specific VLL service endpoint.

Parameters *endpoint-name* — Specify an existing VLL service endpoint name.

force-switchover

Syntax **force-switchover** *sdp-id:vc-id*
no force-switchover

Context tools>perform>service>id

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

Parameters *sdp-id:vc-id* — Specify an existing spoke SDP for the service.

subscriber-mgmt

Syntax **subscriber-mgmt**

Context tools>perform

Description This command enables tools to control subscriber management.

edit-lease-state

Syntax **edit-lease-state sap** *sap-id* **ip** *ip-address* [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*]
edit-lease-state svc-id *service-id* **ip** *ip-address* [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*]

Context **tools>perform>subscr-mgmt**

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

Values *sap-id*:

null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi/vci* |*vpi* |*vpi1.vpi2*]
frame	[*port-id* | *bundle-id*]:*dlci*
cisco-hdlc	*slot/mda/port.channel*

port-id	*slot/mda/port*[.*channel*]	
aps-id	aps-*group-id*[.*channel*]	
	aps	keyword
	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword

	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

ip *ip-address* — Modifies lease state information for the specified IP address.

subscriber *sub-ident-string* — Modifies lease state information for the specified subscriber identification.

sub-profile-string *sub-profile-string* — Modifies lease state information for the specified subscriber profile.

sla-profile-string *sla-profile-string* — Modifies lease state information for the specified SLA profile.

svc-id *service-id* — Modifies lease state information for the specified service ID.

Values 1 — 2147483647

eval-lease-state

Syntax **eval-lease-state** [**svc-id** *service-id*] [**sap** *sap-id*] [**subscriber** *sub-ident-string*] [**ip** *ip-address*]

Context tools>perform>subscr-mgmt

Description This command evaluates lease state information.

Parameters **svc-id *service-id* —** Evaluates lease state information for the specified service.

Values 1 — 2147483647

sap *sap-id* — Evaluates lease state information for the specified SAP.

Values	*sap-id*:	null	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]
		dot1q	[*port-id* | *bundle-id* | *lag-id* | *aps-id*]:*qtag1*
		qinq	[*port-id* | *bundle-id* | *lag-id*]:*qtag1.qtag2*
		atm	[*port-id* | *bundle-id*][:*vpi/vci* | *vpi* | *vpi1.vpi2*]
		frame	[*port-id* | *bundle-id*]:*dlci*
		cisco-hdlc	*slot/mda/port.channel*
		ima-grp	*bundle-id*[:*vpi/vci* | *vpi* | *vpi1.vpi2*]
		port-id	*slot/mda/port*[.*channel*]
		aps-id	aps-*group-id*[.*channel*]
			aps: keyword

	group-id:	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle:	keyword
	type:	ima, ppp
	bundle-num:	1 — 128
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag:	keyword
	id:	1 — 8
	path-id:	a, b
	cc-type:	.sap-net, .net-sap
	cc-id:	0 — 4094
lag-id	lag-*id*	
	lag:	keyword
	id:	1 — 64
qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255]
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

subscriber *sub-ident-string* — Evaluates lease state information for the specified subscriber identification string.

ip *ip-address* — Evaluates lease state information for the specified IP address.

forcerenew

Syntax **forcerenew svc-id** *service-id* {**ip** *ip-address*[/*mask*] | **mac** *ieee-address*}
forcerenew {**interface** *interface-name* | **sap** *sap-id* | **sdp** *sdp-id:vc-id*} [**ip** *ip-address*[/*mask*] |**mac** *ieee-address*]

Context tools>perform>subscr-mgmt

Description This command forces the renewal of lease state.

Parameters **svc-id *service-id* —** Forces renewal of the lease state for the specified service.

Values 1 — 2147483647

sap *sap-id* — Forces renewal of the lease state for the specified SAP.

Values *sap-id*:	null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]
	dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*
	qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*
	atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:*vpi*/*vci* |*vpi* |*vpi1.vpi2*]
	frame	[*port-id* | *bundle-id*]:*dlci*
	cisco-hdlc	*slot/mda/port.channel*
	port-id	*slot/mda/port*[.*channel*]
	aps-id	aps-*group-id*[.*channel*]
	aps	keyword

	group-id	1 — 64
bundle-*type-slot/mda.bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 — 128
bpgrp-id:	**bpgrp**-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 — 1280
ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap]
	cc-id	0 — 4094
lag-id	lag-*id*	
	lag	keyword
	id	1 — 200

qtag1	0 — 4094	
qtag2	*, 0 — 4094	
vpi	NNI	0 — 4095
	UNI	0 — 255
vci	1, 2, 5 — 65535	
dlci	16 — 1022	

ip *ip-address* — Forces renewal of the lease state for the specified IP address.

mac *ieee-address* **—** Forces renewal of the lease state for the specified MAC address.

interface *interface-name* **—** Forces renewal of the lease state for the specified interface name.

re-ident-sub

Syntax **re-ident-sub** *old-sub-ident-string* **to** *new-sub-ident-string*

Context tools>perform>subscr-mgmt

Description This command renames a subscriber identification string.

Parameters *old-sub-ident-string* — Specifies the existing subscriber identification string to be renamed.

new-sub-ident-string — Specifies the new subscriber identification string name.

Tools

This section provides the Tools command reference and hierarchies.

Tools Command Reference

Command Hierarchies

Configuration Commands

```
tools
    — dump
        — lag lag-id lag-id
        — ldp-treetrace {prefix ip-prefix/mask| manual-prefix ip-prefix/mask}[path-destination ip-address] [trace-tree]
        — persistence
            — submgt [record record-key]
            — summary
        — ppp port-id
        — router router-instance
            — dhcp
                — group-if-mapping [clear]
                — group-if-stats [clear]
            — ldp
                — fec prefixip-prefix/mask
                — fec vc-type {ethernet|vlan} vc-id vc-id
                — instance
                — interface [ip-int-name | ip-address]
                — memory-usage
                — peer ip-address
                — session [ip-addr[:label-space] [connection|peer|adjacency]
                — sockets
                — timers
            — mpls
                — ftn [endpoint endpoint | sender sender | nexthop nexthop | lsp-id lsp-id | tunnel-id tunnel-id | label start-label end-label]
                — ilm [endpoint endpoint | sender sender | nexthop nexthop | lsp-id lsp-id | tunnel-id tunnel-id | label start-label end-label]
                — lspinfo
                — memory-usage
            — ospf
            — ospf3
                — abr [detail]
                — asbr [detail]
                — bad-packet interface-name
                — leaked-routes [summary | detail]
```

```
                                    — memory-usage [detail]
                                    — request-list [neighbor ip-address] [detail]
                                    — request-list virtual-neighbor ip-address area-id area-id [detail]
                                    — retransmission-list [neighbor ip-address] [detail]
                                    — retransmission-list virtual-neighbor ip-address area-id area-id [detail]
                                    — route-summary
                                    — route-table [type] [detail]
                                    — te-database [router | link] [summary | detail]
                                — pim
                                    — iom-failures [detail]
                                — rsvp
                                    — psb [endpoint endpoint-address] [sender sender-address] [tunnelid tunnel-id] [lspid lsp-id]
                                    — rsb [endpoint endpoint-address] [sender sender-address] [tunnelid tunnel-id] [lspid lsp-id]
                                    — tcsb[endpoint endpoint-address] [sender sender-address] [tunnelid tunnel-id] [lspid lsp-id]
                                — web-rd
                                    — http-client  [ip-prefix/mask]
                            — service
                                — base-stats [clear]
                                — iom-stats [clear]
                                — l2pt-diags
                                — l2pt-diags clear
                                — l2pt-diags detail
                                — radius-discovery [svc-id service-id]
                                — vpls-fdb-stats [clear]
                                — vpls-mfib-stats [clear]
                            — system-resources slot-number


tools
    — perform
            — aps
                — clear aps-id {protect | working}
                — exercise aps-id {protect | working}
                — force aps-id {protect | working}
                — lockout aps-id
                — request aps-id {protect | working}
            — cron
                — action
                    — stop [action-name] [owner action-owner] [all]
                — tod
                    — re-evaluate
                        — customer customer-id [site customer-site-name]
                        — filter filter-type [filter-id]
                        — service id service-id [sap sap-id]
                        — tod-suite tod-suite-name
            — lag
                — clear-force all-mc
                — clear-force lag-id lag-id [sub-group sub-group-id]
                — clear-force peer-mc ip-address
                — force all-mc {active | standby}
                — force lag-id lag-id [sub-group sub-group-id] {active | standby}
                — force peer-mc peer-ip-address {active | standby}
```

```
        — log
                — test-event
        — router [router-instance]
                — consistency
                — isis
                        — run-manual-spf
                — mpls
                        — cspf to ip-addr [from ip-addr] [bandwidth bandwidth] [include-bitmap
                          bitmap] [exclude-bitmap bitmap] [hop-limit limit] [exclude-address ip-
                          addr [ip-addr ...(up to 8 max)]]
                        — resignal lsp lsp-name path path-name
                        — trap-suppress number-of-traps time-interval
                — ospf [ospf-instance]
                        — refresh-lsas [lsa-type] [area-id]
                        — run-manual-spf externals-only
                — ospf3 [ospf-instance]
                        — refresh-lsas [lsa-type] [area-id]
                        — run-manual-spf externals-only
        — security
                — authentication-server-check server-address ip-address [port port] user-name
                  DHCP client user name password password secret  key  [source-address ip-
                  address] [timeout seconds] [router router-instance]
        — service
                — egress-multicast-group group-name
                        — force-optimize
                — id service-id
                        — endpoint endpoint-name
                                — force-switchover sdp-id:vc-id
                                —  no force-switchover
tools
    — perform
        — subscriber-mgmt
                — edit-lease-state sap sap-id ip ip-address [subscriber sub-ident-string] [sub-profile-
                  string sub-profile-string] [sla-profile-string sla-profile-string]
                — edit-lease-state  svc-id service-id ip ip-address [subscriber sub-ident-string] [sub-
                  profile-string sub-profile-string] [sla-profile-string sla-profile-string]
                — eval-lease-state [svc-id service-id] [sap sap-id] [subscriber sub-ident-string] [ip ip-
                  address]
                — forcerenew svc-id service-id {ip ip-address[/mask]>|mac ieee-address}
                — forcerenew {interface interface-name | sap sap-id|sdp sdp-id:vc-id} [ip ip-
                  address[/mask] |mac ieee-address]
                — re-ident-sub old-sub-ident-string to new-sub-ident-string
```

Standards and Protocol Support

Standards Compliance

IEEE 802.1d Bridging
IEEE 802.1p/Q VLAN Tagging
IEEE 802.1s Multiple Spanning Tree
IEEE 802.1w Rapid Spanning Tree Protocol
IEEE 802.1x Port Based Network Access Control
IEEE 802.3 10BaseT
IEEE 802.3ad Link Aggregation
IEEE 802.3ae 10Gbps Ethernet
IEEE 802.3u 100BaseTX
IEEE 802.3x Flow Control
IEEE 802.3z 1000BaseSX/LX

Protocol Support

OSPF

RFC 1765 OSPF Database Overflow
RFC 2328 OSPF Version 2
RFC 2370 Opaque LSA Support
RFC 3101 OSPF NSSA Option
RFC 3137 OSPF Stub Router Advertisement
RFC 3630 Traffic Engineering (TE) Extensions to OSPF Version 2

BGP

RFC 1397 BGP Default Route Advertisement
RFC 1965 Confederations for BGP
RFC 1997 BGP Communities Attribute
RFC 2385 Protection of BGP Sessions via MD5
RFC 2439 BGP Route Flap Dampening
RFC 2547bis BGP/MPLS VPNs
RFC 2796 BGP Route Reflection: Alternative to Ful-mesh IBGP (previously RFC 1966)
draft-ietf-idr-rfc2796bis-02.txt.
RFC 2858 Multi-protocol Extensions for BGP
draft-ietf-idr-rfc2858bis-09.txt.
RFC 2918 Route Refresh Capability for BGP-4
RFC 3065 Confederations for BGP
draft-ietf-idr-rfc3065bis-05.txt.
RFC 3392 Capabilities Advertisement
RFC 4271 BGP-4 (previously RFC 1771)
RFC 4360 BGP Extended Communities Attribute

IS-IS

RFC 1142 OSI IS-IS Intra-domain Routing Protocol (ISO 10589)
RFC 1195 Use of OSI IS-IS for routing in TCP/IP & dual environments
RFC 2763 Dynamic Hostname Exchange for IS-IS
RFC 2966 Domain-wide Prefix Distribution with Two-Level IS-IS
RFC 2973 IS-IS Mesh Groups
RFC 3373 Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies
RFC 3567 Interfmediate System to Intermediate System (IS-IS) Cryptographic Authentication
RFC 3719 Recommendations for Interoperable Networks using IS-IS
RFC 3784 Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)
RFC 3787 Recommendations for Interoperable IP Networks
draft-ietf-isis-igp-p2p-over-lan-05.txt

LDP

RFC 3036 LDP Specification
RFC 3037 LDP Applicability

IPv6

RFC 1981 Path MTU Discovery for IPv6
RFC 2460 Internet Protocol, Version 6 (IPv6) Specification
RFC 2461 Neighbor Discovery for IPv6
RFC 2462 IPv6 Stateless Address Auto configuration
RFC 2463 Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 Specification

RFC 4644 Transmission of IPv6 Packets over Ethernet Networks
RFC 2529 Transmission of IPv6 over IPv4 Domains without Explicit Tunnels
RFC 2545 Use of BGP-4 Multi-protocol Extension for IPv6 Inter-Domain Routing
RFC 2740 OSPF for IPv6
RFC 3587 IPv6 Global Unicast Address Format
RFC 4007 IPv6 Scoped Address Architecture
RFC 4193 Unique Local IPv6 Unicast Addresses
RFC 4291 IPv6 Addressing Architecture
draft-ietf-ipv6-over-ppp-v2-02
draft-ietf-isis-ipv6-05
draft-ietf-isis-wg-multi-topology-xx.txt

Multicast

RFC 1112 Host Extensions for IP Multicasting (Snooping)
RFC 2236 Internet Group Management Protocol, (Snooping)
RFC 3376 Internet Group Management Protocol, Version 3 (Snooping)
RFC 2362 Protocol Independent Multicast-Sparse Mode (PIM-SM)
RFC 3618 Multicast Source Discovery Protocol (MSDP)
RFC 3446 Anycast Rendevous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)
Draft-ietf-pim-anycast-rp-03
draft-ietf-pim-sm-v2-new-11.txt
draft-ietf-mboned-msdp-mib-01.txt

MPLS

RFC 2702 Requirements for Traffic Engineering over MPLS
RFC 3031 MPLS Architecture
RFC 3032 MPLS Label Stack Encoding
RFC 4379 LSP Ping

RIP

RFC 1058 RIP Version 1
RFC 2082 RIP-2 MD5 Authentication
RFC 2453 RIP Version 2

RSVP-TE

RFC 2430 A Provider Architecture for DiffServ & TE
RFC 3209 Extensions to RSVP for LSP Tunnels
RFC 4090 Fast reroute Extensions to RSVP-TE for LSP Tunnels

DIFFERENTIATED SERVICES

RFC 2474 Definition of the DS Field in the IPv4 and IPv6 Headers
RFC 2597 Assured Forwarding PHB Group
RFC 2598 An Expedited Forwarding PHB
RFC 3140 Per-Hop Behavior Identification Codes

TCP/IP

RFC 768 UDP
RFC 1350 The TFTP Protocol (Rev. 2)
RFC 791 IP
RFC 792 ICMP
RFC 793 TCP
RFC 826 ARP
RFC 854 Telnet
RFC 951 BootP
RFC 1519 CIDR
RFC 1542 Clarifications and Extensions for the Bootstrap Protocol
RFC 1812 Requirements for IPv4 Routers
RFC 2401 Security Architecture for the Internet Protocol
draft-ietf-bfd-mib-00.txtBidirectional Forwarding Detection Management Information Base
draft-ietf-bfd-base-02.txtBidirectional Forwarding Detection
draft-ietf-bfd-v4v6-1hop-02.txtBFD for IPv4 and IPv6 (Single Hop)
draft-bonica-tcp-auth-05.txt, Authentication for TCP-based Routing and Management Protocols

VRRP

RFC 2787 Definitions of Managed Objects for the Virtual Router Redundancy Protocol
RFC 3768 Virtual Router Redundancy Protocol

PPP

RFC 1332 PPP IPCP
RFC 1377 PPP OSINLCP
RFC 1638/2878PPP BCP
RFC 1661 PPP
RFC 1662 PPP in HDLC-like Framing
RFC 1989 PPP Link Quality Monitoring
RFC 2615 PPP over SONET/SDH
RFC 1990 The PPP Multilink Protocol (MP)

ATM

RFC 1626 Default IP MTU for use over ATM AAL5, May 1994
RFC 2514 Definitions of Textual Conventions and OBJECT_IDENTITIES for ATM Management, February 1999
RFC 2515 Definition of Managed Objects for ATM Management, February 1999
RFC 2684 Multiprotocol Encapsulation over ATM Adaptation Layer 5, September 1999
af-tm-0121.000 Traffic Management Specification Version 4.1, March 1999
ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95
ITU-T Recommendation I.432.1 - B-ISDN user-network interface - Physical layer specification: General characteristics
GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996
GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994
AF-ILMi-0065.000 Integrated Local Management Interface (ILMI) Version 4.0
AF-TM-0150.00 Addendum to Traffic Management v4.0 optional minimum desired cell rate indication for UBR

DHCP

RFC 2131 Dynamic Host Configuration Protocol
RFC 3046 DHCP Relay Agent Information Option (Option 82)

RFC 1534 Interoperation between DHCP and BOOTP

VPLS

draft-ietf-l2vpn-vpls-ldp-08.txtVirtual Private LAN Services Using LDP

PSEUDO-WIRE

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 3916 Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)

draft-ietf-pwe3-atm-encap-10.txt

draft-ietf-pwe3-cell-transport-04.txt

draft-ietf-pwe3-ethernet-encap-11.txt

draft-ietf-pwe3-frame-relay-07.txt

draft-ietf-pwe3-control-protocol-17.txt

draft-ietf-l2vpn-vpws-iw-oam-00.txt

draft-ietf-pwe3-vccv-07.txt

draft-ietf-pwe3-oam-msg-map-04-txt

draft-ietf-l2vpn-arp-mediation-04.txt

draft-ietf-pwe3-iana-allocation-15.txt

draft-hart-pwe3-segmented-pw-vccv-01.txt

SONET/SDH

GR-253-CORE SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000

ITU-G.841 Telecommunication Standardization Section of ITU, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum1 issued in July 2002

GR-253-CORE - SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000

RADIUS

RFC 2865 Remote Authentication Dial In User Service

RFC 2866 RADIUS Accounting

SSH

draft-ietf-secsh-architecture.txtSSH Protocol Architecture

draft-ietf-secsh-userauth.txt SSH Authentication Protocol

draft-ietf-secsh-transport.txt SSH Transport Layer Protocol

draft-ietf-secsh-connection.txt SSH Connection Protocol

draft-ietf-secsh- newmodes.txt SSH Transport Layer Encryption Modes

TACACS+

draft-grant-tacacs-02.txt

NETWORK MANAGEMENT

ITU-T X.721: Information technology-OSI-Structure of Management Information

ITU-T X.734: Information technology-OSI-Systems Management: Event Report Management Function

M.3100/3120 Equipment and Connection Models

TMF 509/613 Network Connectivity Model

RFC 1157 SNMPv1

RFC 1657 BGP4-MIB

RFC 1724 RIPv2-MIB

RFC 1850 OSPF-MIB

RFC 1907 SNMPv2-MIB

RFC 2011 IP-MIB

RFC 2012 TCP-MIB

RFC 2013 UDP-MIB

RFC 2096 IP-FORWARD-MIB

RFC 2138 RADIUS

RFC 2206 RSVP-MIB

RFC 2452 IPv6 Management Information Base for the Transmission Control Protocol

RFC 2454 IPv6 Management Information Base for the User Datagram Protocol

RFC 2465 Management Information Base for IPv6: Textual Conventions and General Group

RFC 2558 SONET-MIB

RFC 2571 SNMP-FRAMEWORK-MIB

RFC 2572 SNMP-MPD-MIB

RFC 2573 SNMP-TARGET-&-NOTIFICATION-MIB

RFC 2574 SNMP-USER-BASED-SM-MIB

RFC 2575 SNMP-VIEW-BASED-ACM-MIB

RFC 2576 SNMP-COMMUNITY-MIB

RFC 2665 EtherLike-MIB

RFC 2819 RMON-MIB

RFC 2863 IF-MIB

RFC 2864 INVERTED-STACK-MIB

RFC 2987 VRRP-MIB

RFC 3014 NOTIFICATION-LOG-MIB

RFC 3273 HCRMON-MIB

draft-ietf-disman-alarm-mib-04.txt

draft-ietf-ospf-mib-update-04.txt

draft-ietf-mpls-lsr-mib-06.txt

draft-ietf-mpls-te-mib-04.txt

draft-ietf-mpls-ldp-mib-07.txt

draft-ietf-isis-wg-mib-05.txt

IANA-IFType-MIB

IEEE8023-LAG-MIB

Proprietary MIBs

TIMETRA-APS-MIB.mib

TIMETRA-ATM-MIB.mib

TIMETRA-BGP-MIB.mib

TIMETRA-CAPABILITY-7750-V4v0.mib

TIMETRA-CFLOWD-MIB.mib

TIMETRA-CHASSIS-MIB.mib

TIMETRA-CLEAR-MIB.mib

TIMETRA-FILTER-MIB.mib

TIMETRA-GLOBAL-MIB.mib

TIMETRA-IGMP-MIB.mib

TIMETRA-ISIS-MIB.mib

TIMETRA-LAG-MIB.mib

TIMETRA-LDP-MIB.mib

TIMETRA-LOG-MIB.mib

TIMETRA-MIRROR-MIB.mib

TIMETRA-MPLS-MIB.mib

TIMETRA-NG-BGP-MIB.mib

TIMETRA-OAM-TEST-MIB.mib

TIMETRA-OSPF-MIB.mib

TIMETRA-OSPF-V3-MIB.mib

TIMETRA-PIM-MIB.mib

TIMETRA-PORT-MIB.mib

TIMETRA-PPP-MIB.mib

TIMETRA-QOS-MIB.mib

TIMETRA-RIP-MIB.mib

TIMETRA-ROUTE-POLICY-MIB.mib

TIMETRA-RSVP-MIB.mib

TIMETRA-SECURITY-MIB.mib

TIMETRA-SERV-MIB.mib

TIMETRA-SUBSCRIBER-MGMT-

MIB.mib
TIMETRA-SYSTEM-MIB.mib
TIMETRA-TC-MIB.mib
TIMETRA-VRRP-MIB.mib
TIMETRA-VRTR-MIB.mib

Index

S

T

V