Alcatel-Lucent

Service Access Switch| Release 7.0 Rev.01

7210 SAS-M and 7210 SAS-T OS
Services Guide

3HE09513AAAA

Alcatel·Lucent

Table of Contents

List of Tables

List of Figures

Services Overview

VLL Services

Virtual Private LAN Service

IEEE 802.1ah Provider Backbone Bridging

Internet Enhanced Service

Virtual Private Routed Network Service

Preface

About This Guide

This guide describes service mirroring and Operations, Administration and Management (OAM) and diagnostic tools provided by the 7210 SAS-M, T, platforms and presents examples to configure and implement various tests.

On 7210 SAS devices, not all the CLI commands are supported on all the platforms and in all the modes. In many cases, the CLI commands are mentioned explicitly in this document. In other cases, it is implied and easy to know the CLIs that are not supported on a particular platform.

All the variants of 7210 SAS-M and 7210 SAS-T can be configured in two modes, that is in network mode and in access-uplink mode. In network mode configuration 7210 SAS-M and 7210 SAS-T uses IP/MPLS to provide service transport. In access-uplink mode configuration 7210 SAS-M and 7210 SAS-T uses Ethernet QinQ technology to provide service transport. The mode can be selected by configuring the BOF appropriately.

This guide also presents examples to configure and implement various tests.

Notes:

- This user guide is applicable to all 7210 SAS-M, 7210 SAS-T platforms, unless specified otherwise.
- In either mode, it is expected that the user will only configure the required CLI parameters appropriate for the mode he intends to use. Unless otherwise noted, most of the configuration is similar in both the network mode and Access uplink mode.
- Only 7210 SAS-M and 7210 SAS-T supports access-uplink mode. 7210 SAS-X and 7210 SAS-R6 does not support access-uplink mode.
- On 7210 SAS devices, not all the CLI commands are supported on all the platforms and in all the modes. In many cases, the CLI commands are mentioned explicitly in this document. In other cases, it is implied and easy to know the CLIs that are not supported on a particular platform.

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 7210 SAS M. 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 7210 SAS M, T, X, R6 OS documentation set is composed of the following books:

- 7210 SAS M, T, X, R6 OS Basic System Configuration Guide

  This guide describes basic system configurations and operations.

- 7210 SAS M, T, X, R6 OS System Management Guide

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

- 7210 SAS M, T, X, R6 OS Interface Configuration Guide

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

- 7210 SAS M, T, X, R6 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.

- 7210-SAS-M and 7210 SAS-T OS Services Guide

  This guide describes how to configure service parameters such as customer information and user services.

- 7210 SAS-M, T, X, R6 OS MPLS Guide

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

- 7210 SAS M, T, X, R6 OS OAM and Diagnostic Guide

  This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.

- 7210 SAS-M and 7210 SAS-T OS 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 7210 SAS-series router and related products from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance. If you purchased an Alcatel-Lucent service agreement, contact your welcome center:

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

Getting Started

In This Chapter

This book provides process flow information to configure provision services.

Alcatel-Lucent 7210 SAS 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: Configuration Process

Area	Task	Chapter
Subscribers	Subscriber services	
	Global entities	Configuring Global Service Entities with CLI on page 65
	VLL services	Ethernet Pipe (Epipe) Services on page 144
	VPLS service	Virtual Private LAN Service on page 263
	IES service	Internet Enhanced Service on page 465
	VPRN service	Internet Enhanced Service on page 465
Reference	List of IEEE, IETF, and other proprietary entities.	Standards and Protocol Support on page 1003

Services Command Reference

In This Chapter

This chapter provides the command reference trees for the 7210 SAS services.

Topics include:

- Global Services Commands
- Service Configuration Commands
  - → Cpipe Service Configuration Commands on page 217
  - → on page 218
  - → VPLS Service Configuration Commands on page 382
  - → IES Service Configuration Commands on page 549
  - → VPRN Service Configuration Commands on page 616

SERVICES OVERVIEW

In This Section

This section provides an overview of the 7210 SAS M-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 and an optional service name within a service area. The 7210 SAS-Series service model uses logical service entities to construct a service. In the service model, logical service entities provide a uniform, service-centric configuration, management, and billing model for service provisioning.

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

Note: SDPs are not supported on 7210 SAS-M and 7210 SAS-T devices configured in Access Uplink mode. Only local services can be configured on 7210 SAS-M and 7210 SAS-T configured in access-uplink mode.

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

Service Types

The 7210 SAS-M 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 144.
- Virtual Private LAN Service (VPLS) — A Layer 2 multipoint-to-multipoint VPN. See Virtual Private LAN Service on page 263.
- Internet Enhanced Service (IES) — A routed connectivity service used to provide IP services.This service is supported in 7210 SAS platforms operated in access-uplink mode for only inband management of the node (that is, it cannot be used as for configuring customer service in access-uplink mode). See Internet Enhanced Service on page 465.
- 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 581.

Service Policies

Common to all 7210 SAS-Series connectivity services are policies that are assigned to the service. Policies are defined at a global level and then applied to a service on the router. Policies are used to define 7210 SAS-Series service enhancements. The types of policies that are common to all 7210 SAS-Series connectivity services 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 ingress policy applied to a SAP specifies the number of meters, meter characteristics (such as forwarding class, committed, and peak information rates, etc.) and the mapping of traffic to a forwarding class. A QoS egress policy defines the queue characteristics (such as CBS, CIR, PIR). A QoS policy must be created before it can be applied to a SAP. A single ingress and egress QoS policy can be associated with a SAP. A single access egress QoS policy can be associated with a port.

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

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

- Scheduler policies define the operating parameters (such as scheduling algorithm, weights per priority). They are associated with physical ports.

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

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

Alcatel-Lucent Service Model

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

The 7210 SAS devices configured in access-uplink mode supports QinQ/Dot1q Layer 2 uplinks to transport the services to the provider edge in a hierarchical configuration, whereas 7210 SAS devices configured in network mode support MPLS uplinks to transport the services.

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

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

Service provisioning uses logical entities to provision a service where additional properties can be configured for bandwidth provisioning, QoS, security filtering, 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 29)
- Service Access Points (SAPs) (see page 29)
- Service Distribution Points (SDPs) (see page 41) (for distributed services only)

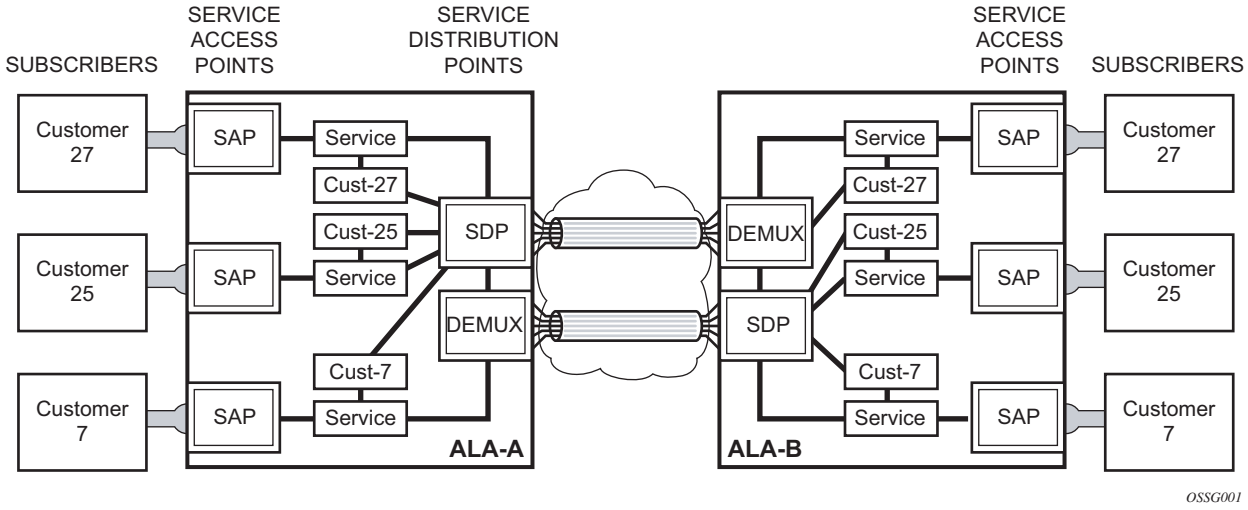

Figure 1: Service Entities for 7210 SAS devices configured in Network Mode

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 7210 SAS M-Series router (Figure 3). The SAP configuration requires that slot, MDA, and port information be specified. The slot, MDA, and port parameters must be configured prior to provisioning a service (see the Cards, MDAs, and Ports sections of the 7210 SAS OS Interface Configuration Guide).

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

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

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

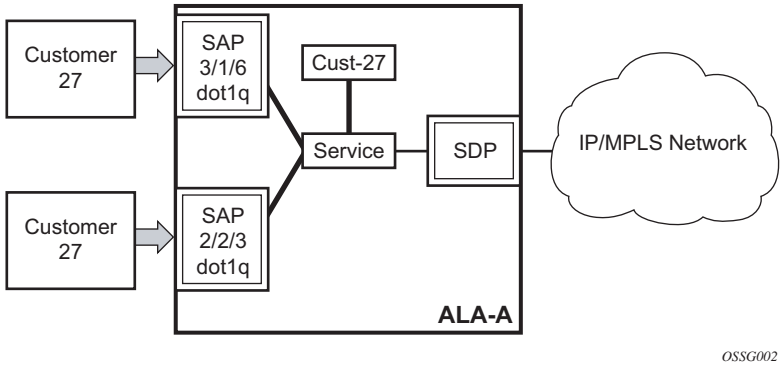

Figure 2: Service Access Point (SAP) for 7210 SAS configured in Network Mode

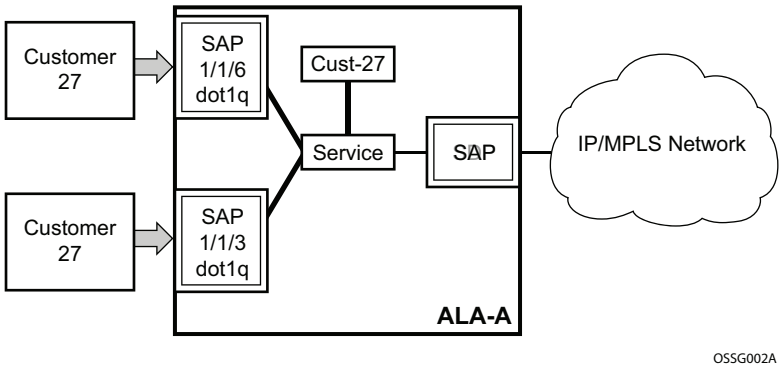

Figure 3: Multiple SAPs in a service using QinQ uplinks in 7210 SAS configured in access-uplink mode

SAP Encapsulation Types and Identifiers

The encapsulation type is an access property of a service Ethernet port. The appropriate encapsulation type for the port depends on the requirements to support multiple services on a single port on the associated SAP and the capabilities of the downstream equipment connected to the port. For example, a port can be tagged with IEEE 802.1Q (referred to as dot1q) encapsulation in which each individual tag can be identified with a service. A SAP is created on a given port 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 4).. The encapsulation ID used to distinguish an individual service is the VLAN ID in the IEEE 802.1Q header. For example, the port is connected to a Ethernet switch (for example, a 7210 SAS E) with multiple downstream customers.
- QinQ — The QinQ encapsulation type adds a IEEE 802.1Q tag to the 802.1Q tagged packets entering the network to expand the VLAN space by tagging tagged packets, producing a double tagged frame. 7210 SAS M OS supports QinQ encapsulation for

access ports in network mode. In access-uplink mode, QinQ encapsulation is supported for both access port and access uplink ports.

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

- QinQ — The QinQ encapsulation type adds a IEEE 802.1Q tag to the 802.1Q tagged packets entering the network to expand the VLAN space by tagging tagged packets, producing a double tagged frame. On the 7210 SAS E, QinQ encapsulation is supported only on access uplink ports.

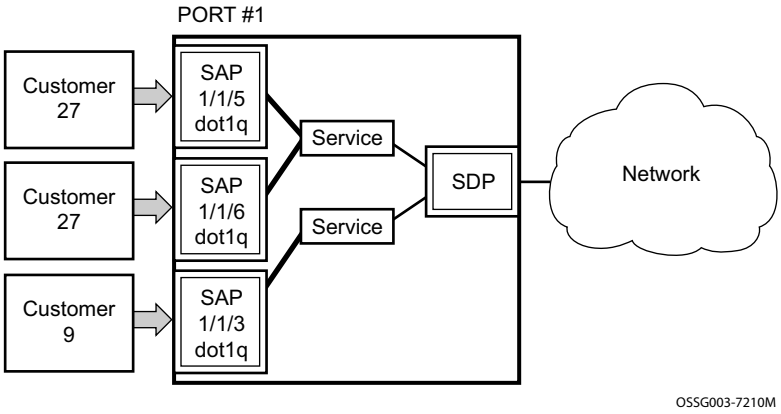

Figure 4: Multiple SAPs on a Single Port (7210 in Network Mode)

Services and SAP Encapsulations

Table 2 lists the service and SAP Encapsulation information for Ethernet ports:

Table 2: Service and Encapsulation

Port Type	Encapsulation
Ethernet	Null
Ethernet	Dot1q
Ethernet	QinQ

Default SAP on a Dot1q Port

This feature introduces default SAP functionality on Dot1q-encapsulated ports. On a dot1q-encapsulated port where a default SAP is configured, all packets with q-tags not matching any explicitly defined SAPs will be assigned to this SAP. SAPs with default Dot1q encapsulation are supported in VPLS and Epipe services. Dot1q Default SAP are not supported in VPRNs. In this context, the character "*" indicates default which means allow through. The default SAP also accepts untagged or priority tagged packets. A default SAP must be configured explicitly. When a default SAP is not configured explicitly, packets not matching any explicitly defined SAPs will be dropped.

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

In this type of environment, logically two SAPs exist, a management SAP and a service SAP. The management SAP can be created by specifying a VLAN tag 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.
- 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.
- 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.

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

Default QinQ SAPs (notation - *.*) are used in ring ports to avoid the need to configure services on all the intermediate nodes in the ring which are transiting the service. Default QinQ SAPs matches all VLAN tagged traffic which is not classified into any other SAP configured on the same port. Only one EPIPE service with default QinQ SAPs is needed for transit service traffic on access-uplink ports. Default QinQ SAPs are allowed only on access-uplink ports and access ports. It can co-exist with 0.* SAP on an access-uplink or access port. A default QinQ SAP accepts only tagged packets. Untagged packets or priority tagged packets are not accepted on Default QinQ SAPs.

When an EPIPE service With default QinQ SAPs on the ring ports is used for transit traffic in a ring deployment, no protection mechanism (example: STP or G.8032) is supported for Default QinQ SAPs. The upstream or head-end node on which the service originates must ensure the correct path on the ring is selected using either G.8032 or STP. When a VPLS service with default QinQ SAPs on the ring ports is used for transit traffic in a ring deployment, users can use either G8032 or M-VPLS with xSTP for ring protection. When using G8032, the state of the default QinQ SAPs in the VPLS service can be managed using a separate G8032 control instance.
NOTE: G8032 control instance cannot use Default QinQ SAPs.

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

Following features are available for use with Default QinQ SAPs configured in EPIPE and VPLS service (unless explicitly specified, below listed features are applicable for both EPIPE and VPLS service):

For Default QinQ SAPs on either access ports or access-uplink ports:

- MAC learning and aging is available for use in a VPLS service
- Per SAP MAC limit is available for use in a VPLS service
- Mac-move detection and Mac-pinning is available for use in a VPLS service
- Discard-unknown and discard-unknown-source is available for use in a VPLS service
- ETH-CFM and Y.1731 is not available for use
- STP (and all its different flavors) cannot be enabled in the service with Default QinQ SAPs
- MVPLS with xSTP can be used for loop prevention. The Default QinQ SAPs inherit the state from the associated MVPLS instance.
- G.8032 control instance cannot be configured in a service with Default QinQ SAP
- G8032 can be used for loop prevention in ring deployments, where the Default QinQ SAPs are configured on the ring ports in a VPLS service. A separate G8032 control instances needs to be configured for use on the ring ports and the service with Default QinQ ports needs to be associated with this G8032 control instance
- IGMP snooping is not available for use

- L2PT and BPDU translation is not available for use
- IP interface in a VPLS service is not supported in a service using this SAP

For Default QinQ SAPs created on Access-uplink Port:

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

For Default QinQ SAPs created on access ports:

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

Configuration Notes for use of Default QinQ SAPs for transit service in a ring deployment

- If an Epipe service is used with Default QinQ SAPs on the ring ports for transit service in a ring deployment, no protection mechanism is available for the transit service (that is, Epipe service with the Default QinQ SAPs on ring ports). Both Epipe and VPLS services which are originating on different nodes in the ring can use the transit service. Protection/ Loop-detection mechanisms can be implemented for VPLS service configured in the ring nodes, by using MVPLS with xSTP on the nodes where the VPLS service is configured. No protection mechanisms are available for use with Epipe services on the node that originates the service.
- If a VPLS service is used with Default QinQ SAPs on the ring ports for transit service in a ring deployment, either MVPLS/xSTP or G8032 can be used to protect the transit service (that is, VPLS service with the Default QinQ SAPs on ring ports). In this case, VPLS service which are originating on different nodes in the ring and use the transit VPLS service are also protected. Epipe services which are originating on different nodes in the ring cannot use the transit VPLS service.
- When using VPLS service with Default QinQ SAPs for transit service with either G8032 or MVPLS with xSTP configured for protection, load-balancing of the traffic based on the VLAN IDs is not possible. If load-balancing is desired then it is better to use Epipe service with Default QinQ SAPs as the transit service.

SAP Configuration Considerations (applicable for both Network mode and access-uplink mode)

When configuring a SAP, consider the following (applicable to both network mode and access-uplink mode):

- A SAP is a local entity and only locally unique to a given device. The same SAP ID value can be used on another 7210 SAS-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.
- A SAP is owned by and associated with the service in which it is created in each router.
- A port with a dot1q encapsulation type means the traffic for the SAP is identified based on a specific IEEE 802.1Q VLAN ID value. The VLAN ID is stripped off at SAP ingress and the appropriate VLAN ID placed on at SAP egress. As a result, VLAN IDs only have local significance, so the VLAN IDs for the SAPs for a service need not be the same at each SAP.

- If a port is administratively shutdown, all SAPs on that port will be operationally out of service.
- QinQ access SAPs of type Q1.0 is not supported.
- A SAP cannot be deleted until it has been administratively disabled (shutdown).
- Each SAP can have one each of the following policies assigned:
  - → Ingress filter policy
  - → Egress filter policy
  - → Ingress QoS policy
  - → Accounting policy

  Note: Access-egress QoS policy is assigned per access port.

- SAPs using connection-profile (to specify dot1q VLAN ranges) can be configured in a service only when svc-sap-type is set to 'dot1q-range'.
- When a service is configured to use svc-sap-type 'dot1q-range', the outermost V-LAN tag of the packets are not stripped when the packet is received on access port ingress. For more information on processing behavior for this type of service, see “Ethernet Pipe (Epipe) Services on page 144 ” section.

QinQ SAP Configuration restrictions for 7210 SAS in Network mode only

Listed below are the QinQ access SAP configuration guidelines for 7210 SAS in Network mode only.

The guidelines listed below are not applicable when the 7210 SAS- M devices are configured in access uplink mode and access uplink SAPs are in use.

- Processing of tagged packets received on SAPs configured in a service in which a QinQ SAP is also in use (not applicable when a QinQ SAP is not provisioned in a service).
- When a QinQ SAP is configured in a service, the number of VLAN tags in the packets received on NULL SAP, Dot1q SAP and QinQ SAP configured in the same service should match the number of VLAN tags implied by the port encapsulation mode. Packets that do not match are dropped by the hardware. That is, packets received with more than two VLAN tags on a QinQ SAP are dropped, packets received with more than one VLAN tag on a Dot1q SAP are dropped and packets received with tags (even packet with a priority tag) on a NULL SAP are dropped. Henceforth in this document, such packets are referred to as extra-tag packets.

- When a QinQ SAP is configured in a service, the number of VLAN tags in the packets received on the VC/pseudowire of type ‘vc-vlan’ should be exactly one and packets received on the VC/pseudowire of type ‘vc-ether’ should contain no tags (not even priority tag). If either case, packets that contain more number of VLAN tags than the number mentioned above are dropped. Henceforth the document refers to such packets as extra-tag packets.
- The system will provide a limited amount of counters to count the number of extra-tag packets dropped on SAP ingress. These counters are intended for diagnostic use.
- Table 3 displays the SAP types allowed in a service when QinQ SAP is in use:

Table 3: SAP types in a service when QinQ SAP is in use (Network mode operation)

SAP configured in the service	SAPs Not Allowed for configuration in the same service
QinQ	Q.* SAP, Dot1q Default SAP
Q.*	Q1.Q2
Dotq1 default SAP	Q1.Q2

0.* QinQ SAP configured in the service will accept only untagged or priority tagged packets, irrespective of whether a QinQ SAP is configured in the service or not.

NOTE: 7210 supports a mechanism to transport QinQ packets in an Epipe with 2 or more tags, with some restrictions. For more information, see “Epipe chapter” .

SAP configuration notes when operating the 7210 SAS devices in Access-Uplink mode only

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

The Table 4 provides details of SAP and service combinations allowed in access-uplink mode

Table 4: SAP and Service Combinations for 7210 SAS-M and 7210 SAS-T in access-uplink

svc-sap-type	Access SAPs	Access Uplink SAPs
null-star	Null SAP,dot1q Default SAP, Default QinQ SAP (*.* SAP)	Q.* SAP, Default QinQ SAP (*.* SAP)

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

mode

- 'svc-sap-type' parameter value determines the type of SAPs that are allowed to be provisioned in a service.
- A physical port can have only one SAP to be part of one service. Multiple SAPS can be defined over a physical port but each of these SAPs should belong to a different service.
- In the case of a service's sap-type is specified as **dot1q-preserve**, all the SAPs configured in the service must have the same VLAN ID. The outermost VLAN tag of the packets received on access port is not stripped, when svc-sap-type is set to dot1q-preserve.
- Dot1q Default SAP cannot be configured when svc-sap-type is set to 'any'
- When svc-sap-type is set to 'any' for a NULL SAP, the system processes and forwards only packets with no VLAN tag (that is, untagged). All other packets with one or more VLAN tags (even those with priority tag only) are not processed and dropped. Users can use the service with svc-sap-type set to 'null-star', to process and forward packets with one or more tags (including priority tag) on a null SAP.
- An ingress QoS policy and accounting policy is assigned per access uplink port and cannot be assigned per access uplink SAP.
- The **Default QinQ** SAP processes only tagged packets received on a QinQ port. All tagged packets that do not match the specific SAP tags configured on the same port are processed by this SAP. The **Default QinQ** SAP cannot process un-tagged packets, even if **0.*** SAP is not configured for use on that port.

The Default QinQ SAPs is available for use with 0.* SAPs configured on the same port or in the same service. It is available for use with another default QinQ SAP configured in the same service

(on a different port).In a VPLS service, the Default QinQ SAP is available for use with any other SAP type configured in a service configured with svc-sap-type parameter set to "null-star".

- → SAPs using connection-profile (to specify dot1q VLAN ranges or individual VLAN IDs) can be configured in a service only when svc-sap-type is set to 'dot1q-range'.
- → When a service is configured to use svc-sap-type 'dot1q-range', the outermost V-LAN tag of the packets are not stripped when the packet is received on access port ingress. For more information, see “Ethernet Pipe (Epipe) Services on page 144” chapter for processing behavior for this type of service.

Service Distribution Points (SDPs)

Note: SDPs are not supported by 7210 SAS devices configured in Access Uplink mode.

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

An SDP has the following characteristics:

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

An SDP from the local device to a far-end router requires a return path SDP from the far-end 7210 SAS-Series back to the local router. Each device must have an SDP defined for every remote router to which it 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) 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 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.

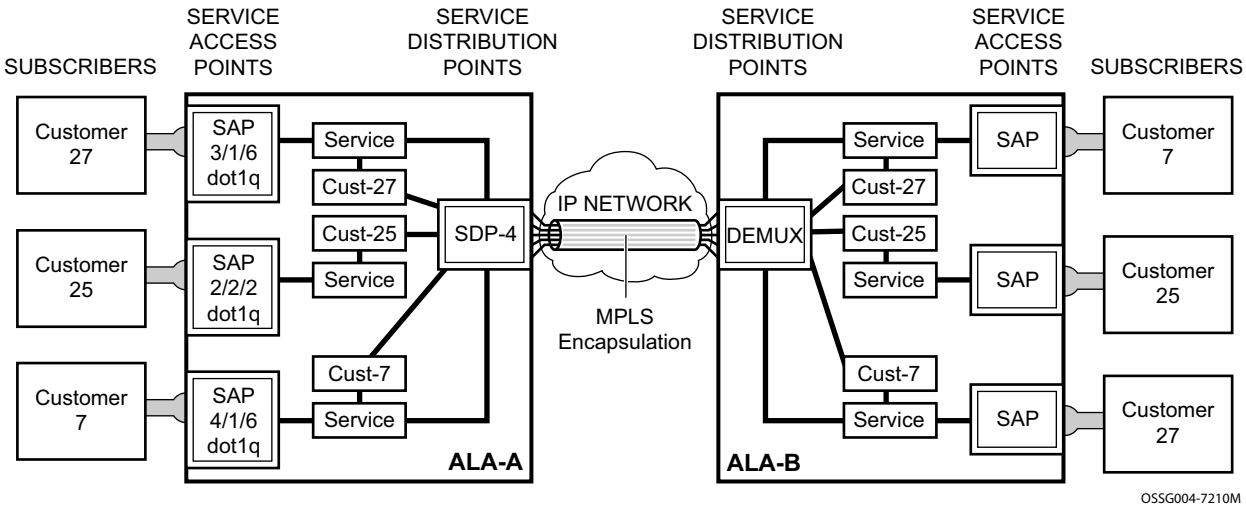

Figure 5: MPLS Service Distribution Point (SDP) Pointing From ALA-A to ALA-B

Spoke and MESH SDPs

NOTE: The 7210 SAS-M and 7210 SAS-T network mode supports SDP. The 7210 SAS-M and 7210 SAS-T in access-uplink mode do not support SDPs.

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

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

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

SDP Using BGP Route Tunnel

SDP is enhanced to use BGP route tunnel to extend inter-AS support for L2 and L3 VPN services. An SDP can be configured to use MPLS transport method. MPLS SDP support is enhanced to allow a BGP route tunnel to reach the far-end PE. A single method of tunneling is allowed per SDP (for example, LDP, RSVP-TE LSP or BGP route tunnel). BGP route tunnel method is excluded if multi-mode transport is enabled for an SDP.

For inter-AS far-end PE, next-hop for BGP route tunnel must be one of the local ASBR. The LSP type selected to reach the local ASBR (BGP labeled route next-hop) must be configured under the BGP global context. LDP must be supported to provide transport LSP to reach the BGP route tunnel next-hop.

Only BGP route labels can be used to transition from ASBR to the next-hop ASBR. The global BGP route tunnel transport configuration option must be entered to select an LSP to reach the PE node from ASBR node. On the last BGP segment, both “BGP+LDP” and LDP routes may be available to reach the far-end PE from the ASBR node. LDP LSP must be preferred due to higher protocol priority. This leads to just one label besides other labels in stack to identify VC/VPN at far-end PE nodes.

SDP Keepalives

SDP keepalives actively monitor 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:

- Admin up/admin down state
- Hello time
- Message length
- Max drop count
- Hold down time

SDP keepalive echo request messages are only sent when the SDP is completely configured and administratively up and SDP keepalives 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 71.

Mixed-LSP Mode of Operation

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

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

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

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

In the case of the RSVP/LDP SDP, the service manager will program the NHLFE(s) for the active LSP type preferring the RSVP LSP type over the LDP LSP type. If no RSVP LSP is configured or all configured RSVP LSPs go down, the service manager will re-program the linecard with the LDP LSP if available. If not, the SDP goes operationally down.

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

Note however, that LDP uses a tunnel down damp timer which is set to three seconds by default. When the LDP LSP fails, the SDP will revert to the RSVP LSP type after the expiry of this timer. For an immediate switchover this timer must be set to zero.
Use the **configure>router>ldp>tunnel-down-damp-time** command. For more information, see 7210 SAS M, X, T, and R6 OS MPLS User Guide.

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

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

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

G.8032 Ethernet Ring Protection Switching

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

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

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

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

Overview of G.8032 Operation

R-APS messages that carry the G.8032 protocol are sent on dedicated protocol VLAN called ERP VLAN (or Ring Control Instance). In a revertive case, G.8032 Protocol ensures that one Ring Protection Link (RPL) owner blocks the RPL link. R-APS messages are periodically sent around in both directions to inform other nodes in the Ring about the blocked port in the RPL owner node. In non-revertive mode any link may be the RPL link.Y.1731 Ethernet OAM CC is the basis of the RAPs messages. Y.1731 CC messages are typically used by nodes in the ring to monitor the health of each link in the ring in both directions. However CC messages are not mandatory. Other link layer mechanisms could be considered – for example LOS (Loss of Signal) when the nodes are directly connected.

Initially each Ring Node blocks one of its links and notifies other nodes in the ring about the blocked link. Once a ring node in the ring learns that another link is blocked, the node unblocks its blocked link possibly causing FDB flush in all links of the ring for the affected service VLANs, controlled by the ring control instance. This procedure results in unblocking all links but the one link and the ring normal (or idle) state is reached. In revertive mode the RPL link will be the link that is blocked when all links are operable after the revert time. In non-revertive mode the RPL link is no different that other ring links. Revertive mode offers predictability particularly when there are multiple ring instances and the operator can control which links are block on the different instances. Each time there is a topology change that affects Reachability, the nodes may flush the FDB and MAC learning takes place for the affected service VLANs, allowing forwarding of packets to continue. Figure 6 depicts this operational state:

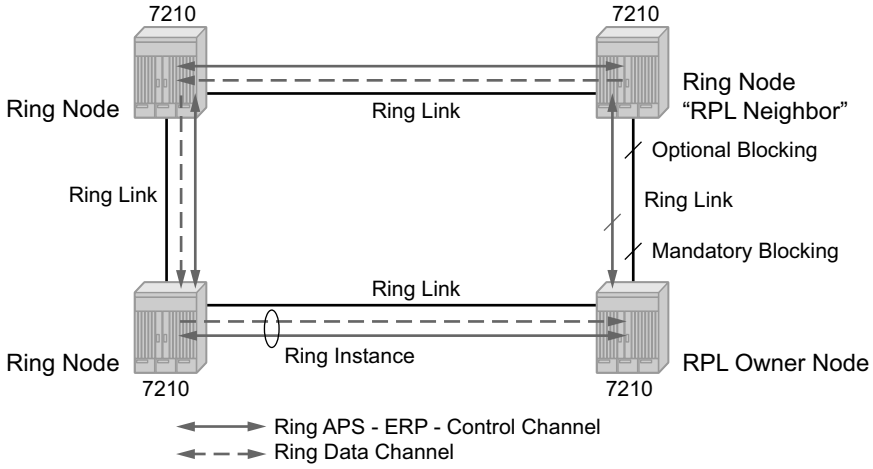

Figure 6: G.8032 Ring in the Initial State

When a ring failure occurs, a node or nodes detecting the failure (enabled by Y.1731 OAM CC monitoring) send R-APS message in both directions. This allows the nodes at both ends of the failed link to block forwarding to the failed link preventing it from becoming active. In revertive mode, the RPL Owner then unblocks the previously blocked RPL and triggers FDB flush for all

nodes for the affected service instances. The ring is now in protecting state and full ring connectivity is restored. MAC learning takes place to allow Layer 2 packet forwarding on a ring. The following picture depicts the failed link scenario.

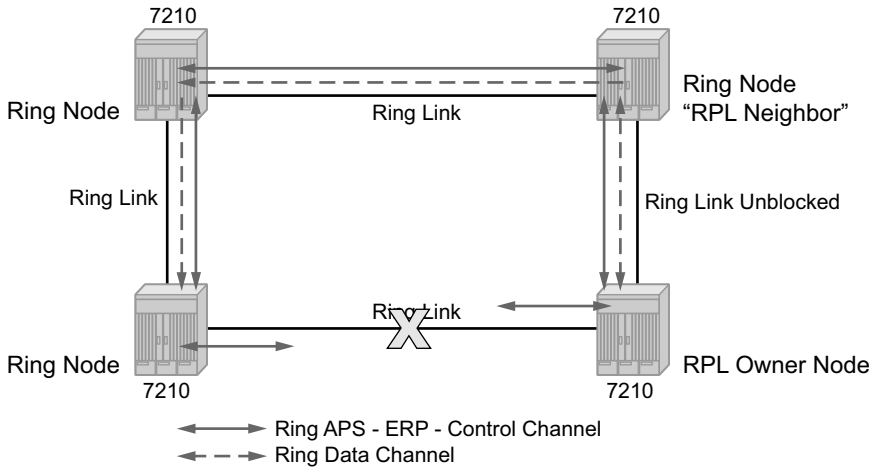

Figure 7: 0-1 G.8032 Ring in the Protecting State

Once the failed link recovers, the nodes that blocked the link again send the R-APS messages indicating no failure this time. This in turn triggers RPL Owner to block the RPL link and indicate the Blocked RPL link the ring in R-APS message, which when received by the nodes at the recovered link cause them to unblock that link and restore connectivity (again all nodes in the ring perform FDB Flush and MAC learning takes place). The ring is back in the normal (or idle) state.

Within each path, Y.1731 Maintenance Entity Group (MEG) Endpoints (MEPs) are used to exchange R-APS specific information (specifically to co-ordinate switchovers) as well as optionally fast Continuity Check Messages (CCM) providing an inherent fault detection mechanism as part of the protocol. Failure detection of a ring path by one of the mechanisms triggers to activate the protection links. Upon failure, re-convergence times are a dependent on the failure detection mechanisms. In the case of Y.1731, the CCM transmit interval determines the response time. The 7210 SAS device supports 100ms (millisecond) message timers that allows for quicker restoration times. Alternatively, 802.3ah (Ethernet in the First Mile) or simple Loss of Signal can act as a trigger for a protection switch where appropriate. In case of direct connectivity between the nodes, there is no need to use Ethernet CC messaging for liveliness detection.

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

G.8032 supports multiple data channels (VIDs) or instances per ring control instance (R-APS tag). G.8032 also supports multiple control instances such that each instance can support RPLs on different links providing for a load balancing capability however once services have been assigned to one instance the rest of the services that need to be interconnected to those services must be on

the same instance. In other words each data instance is a separate data VLAN on the same physical topology. When there is any one link failure or any one node failure in the ring, G.8032 protocols are capable of restoring traffic between all remaining nodes in these data instances.

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

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

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

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

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

The Figure 8 below illustrates a resilient Ring Service. In the ring example, a PBB ring (solid line) using VID 500 carries 2 service VLANs on I-SID 1000 and 1001 for Service VIDs (Dot1q 100 and QinQ 400.1 respectively). The RPL for the PBB ring is between A and B where B is the RPL owner. Also, illustrated in the figure below is a QinQ service on the (dotted line) ring that uses Dot1q VID 600 for the ring to connect service VLAN 100.50. The two rings have RPLs on different nodes which allow a form of load balancing. The example serves to illustrate that service encapsulations and ring encapsulation can be mixed in various combinations. Also, note that neither of the rings is a closed loop. A ring can restore connectivity when any one node or link fails to all remaining nodes within the 50ms transfer time (signaling time after detection).

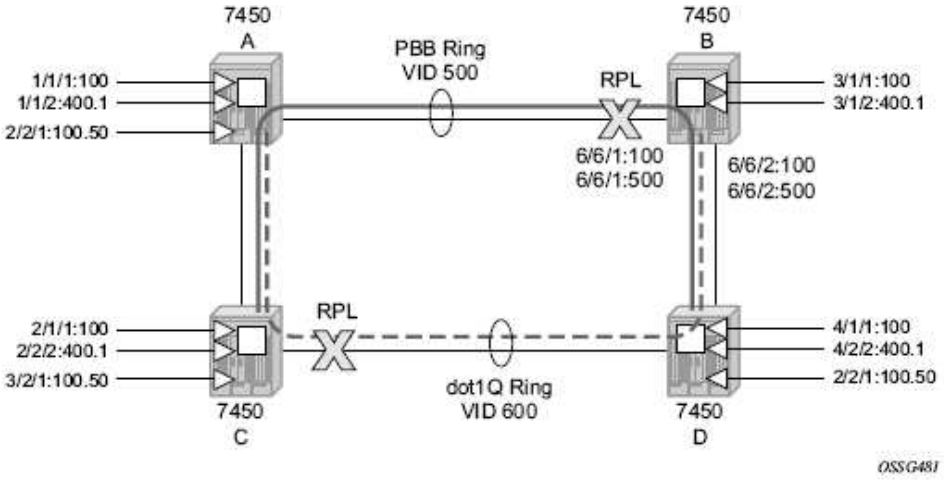

Figure 8: 0-3 Ring Example

Sample Configuration:

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

service
     vpls 10 customer 1 create // Ring APS SAPs
```

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

service
     epipe 100 pbb-epipe   // CPE traffic
          description " PBB epipe service for CPE"
                 pbb-tunnel 40 backbone-dest-mac 00:bb:bb:bb:bb:bb isid 100
            sap 3/1/1:100 create
                      description "Default sap description for service id 100"
            exit
                 no shutdown
     exit
```

Ethernet Ring Sub-Rings

Ethernet Sub-Rings offer a dual redundant way to interconnect rings. The 7210 SAS supports Sub-Rings connected to major rings and a sub-ring connected to a VPLS (LDP based) for access rings support in VPLS networks. Figure 9 illustrates a Major ring and Sub Ring scenario. In this scenario, any link can fail in either ring (ERP1 or ERP2) and each ring is protected. Furthermore, the sub ring (ERP2) relies on the major Ring (ERP1) as part of its protection for the traffic from C and D. The nodes C and D are configured as inter connection nodes.

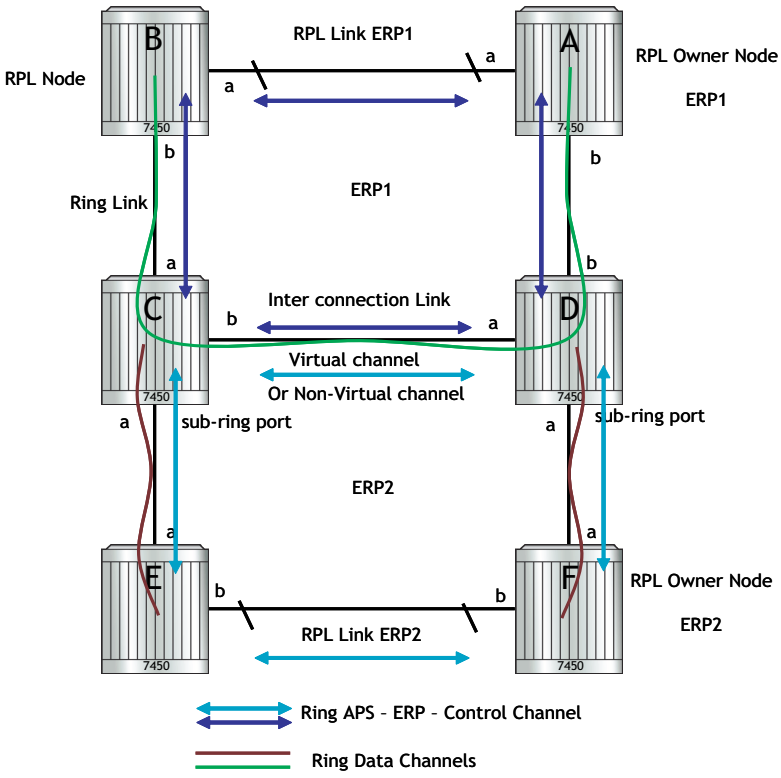

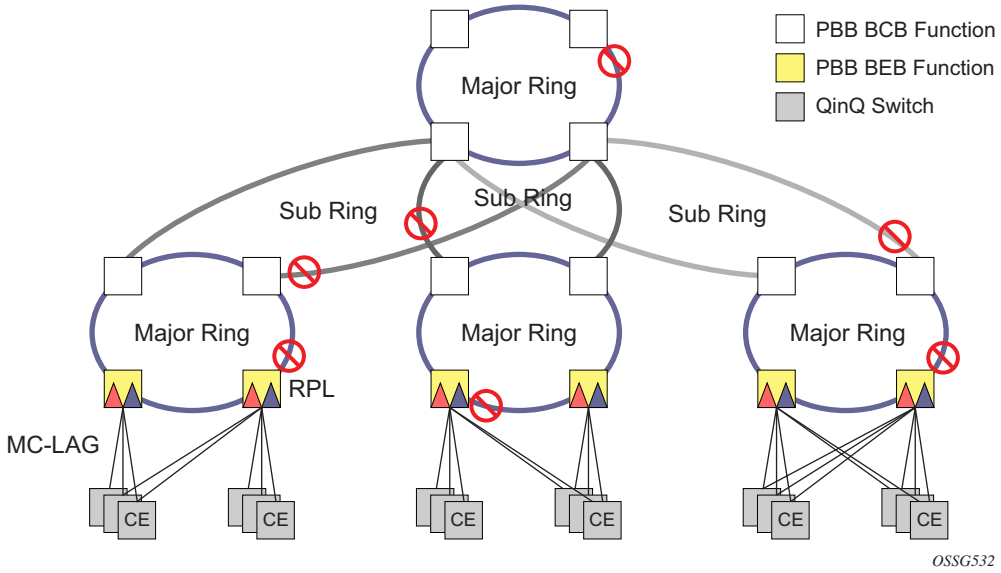

Figure 9: 0-4 G.8032 Sub-Ring

Sub-Rings and Major Rings run similar state machines for the ring logic, however there are some differences. When Sub-Rings protect a link, the flush messages are propagated to the major ring. (A special configuration allows control of this option on the 7210 SAS.) When major rings change topology, the flush is propagated around the major ring and does not continue to any sub-rings. The reason for this is that Major Rings are completely connected but Sub-Rings are dependent on another ring or network for full connectivity. The topology changes need to be propagated to the other ring or network usually. Sub-Rings offer the same capabilities as major rings in terms of control and data so that all link resource may be utilized.

Virtual and Non-Virtual Channel

The following illustrates a sample Sub-Ring using virtual-link configuration on Node C, interconnecting node:

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

    exit

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

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

# Data configuration for the Sub-Ring

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

```
        description "Local Data SAP"
            stp shutdown
        no shutdown
    exit

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

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

Ethernet Ring Sub Ring using non-virtual-link

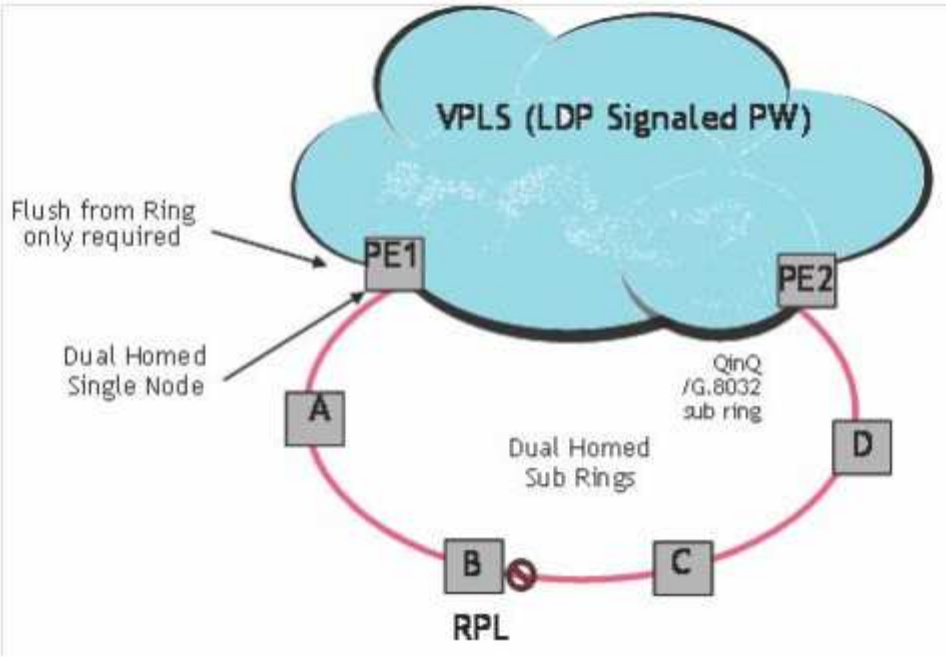

Figure 10: 0-6 Sub-Ring Homed to VPLS

The following illustrates a sample Sub-Ring using non-virtual-link configuration on PE1, interconnecting node:

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

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

All the Sub Ring nodes part of Sub Ring with non-virtual-link should be configured with “sub-ring non-virtual-link” option.

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

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

```
      mesh-sdp 5006:5 create
      exit

      no shutdown
  exit
# Control Channel for Sub-Ring using non-virtual-link on sub-Ring nodes:
vpls 1 customer 1 create
            stp
                shutdown
            exit
            sap 1/1/1:1.1 eth-ring 1 create
                stp
                    shutdown
                exit
            exit
            sap 1/1/2:1.1 eth-ring 1 create
                stp
                    shutdown
                exit
            exit
            no shutdown
        exit
```

Lag Support

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

OAM Considerations

Ethernet CFM can be enabled on each individual path under an Ethernet ring. Only down MEPs can be configured on each of them and CCM sessions can be enabled to monitor the liveliness of the path using interval of 100 msec. Different CCM intervals can be supported on the path a and path b in an Ethernet ring. CFM is optional if hardware supports Loss of Signal for example.

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

QoS Considerations

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

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

the NC traffic should be avoided. The operator must configure different QoS Policies to avoid congestion for the CCM forwarding class by controlling the amount of traffic assigned into the corresponding queue.

Details of the Ethernet ring applicability in the services solution can be found in the respective Layer 2 sections of the 7210 SAS-M OS Services Guide.

Support Service and Solution Combinations

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

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

Configuration guidelines for G.8032

The following are the configuration guidelines for G.8032:

- For 7210 SAS-M devices in network mode, to improve service fail-over time due to failures in the ring path, users can use the CLI command “config> system>resourceprofile>g8032-fast-flood-enable”. When fast flood is enabled, on a failure detection in one of the paths of the eth-ring, along with MAC flush the system starts to flood the traffic onto the available path. The resources needed for this functionality are shared with filters and affects filter scaling. For more information refer to the command description of the command g8032-fast-flood-enable in the 7210 SAS-M,X, T, and R6 Interface configuration guide for more details. For 7210 SAS-M and 7210 SAS-T devices in access-uplink mode, to improve the service fail-over time due to failures in the ring path, fast flood is enabled by default. On a failure detection in one of the paths of the eth-ring, along with MAC flush the system starts to flood the traffic onto the available path. No explicit user configuration is needed for this and it does not affect scaling for filters.Down MEPs used with services and G.8032 share common hardware resources.
- Service level MEPs are not available on all SAPs tied to an eth-ring instance on a port.
- G8032 instances cannot be configured over a LAG.

Service Creation Process Overview

Figure 11 displays the overall process to provision core and subscriber services.

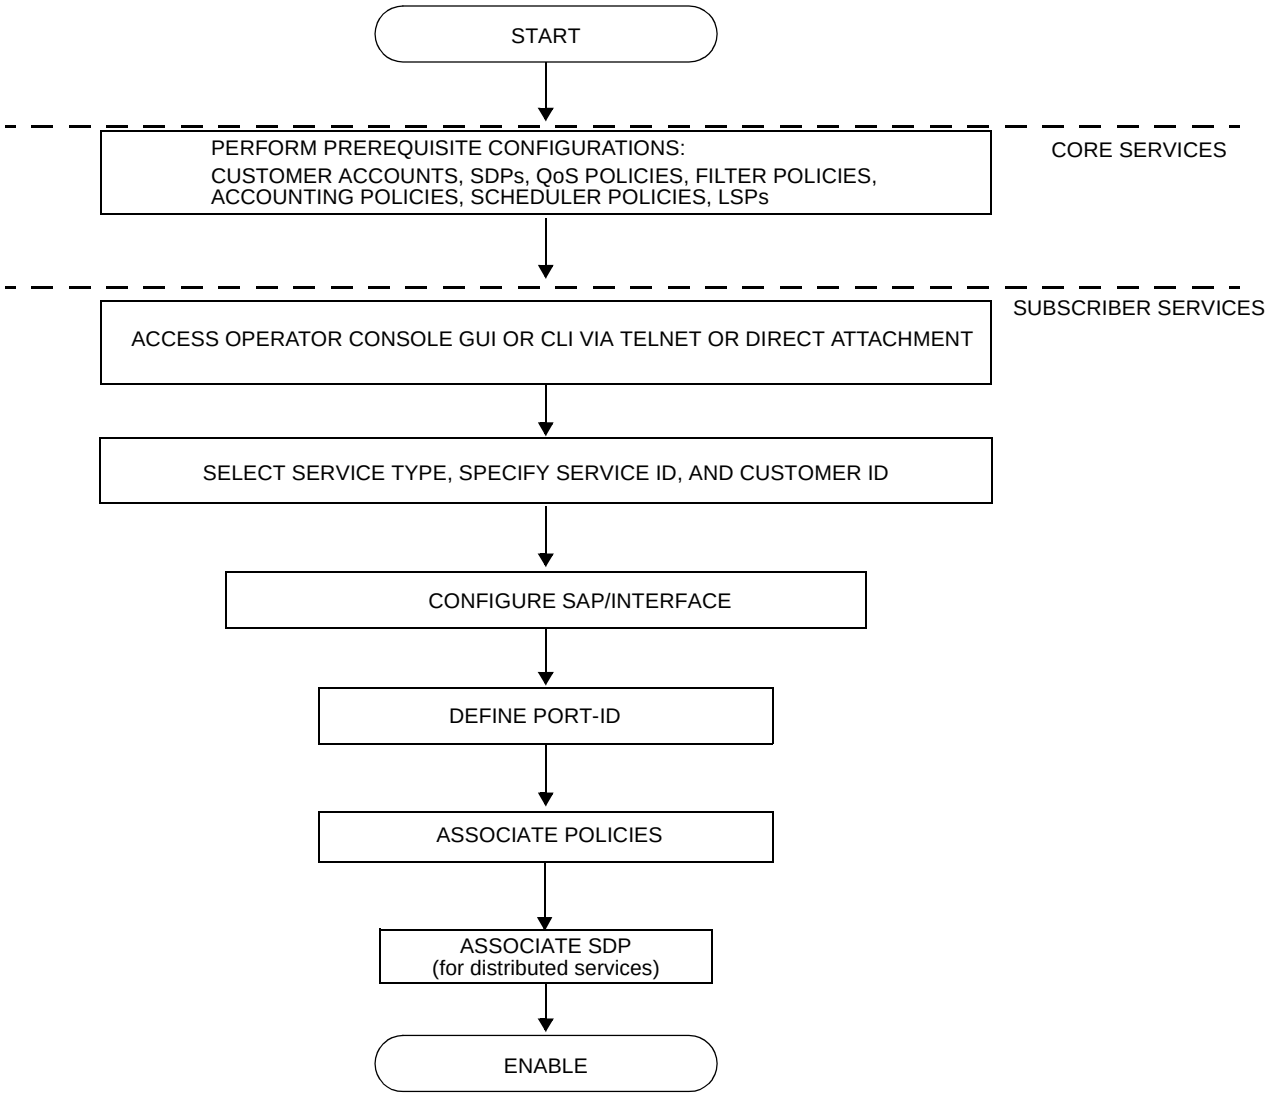

Figure 11: Service Creation and Implementation Flow

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).

Phase 2: Service Administration

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

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

Phase 3: Service Provisioning

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

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 SDPs (Not applicable for devices configured in Access Uplink mode)

Subscriber services tasks include the following:

- Create Epipe and VPLS services.
- Create a VPRN service (Supported only when operating in Network mode)
- Bind SDPs (Not applicable for 7210 SAS devices configured in Access Uplink mode)
- Configure interfaces (where required) and SAPs
- Create exclusive QoS and filter policies

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

Configuring Global Service Entities with CLI

This section provides information to create subscriber (customer) accounts 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 69
- Services:
  - → Ethernet Pipe (Epipe) Services on page 144
  - → VPLS on page 329
  - → IES on page 539
- Service Access Points (SAPs)
  - → Ethernet Pipe (Epipe) Services on page 144
  - → VPLS SAP on page 342

Basic Configuration

The most basic service configuration must have the following:

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

- An associated SDP for distributed services in the network mode. The SDPs are not supported on the access uplink mode.

The following example provides an Epipe service configuration displaying the SDP and Epipe service entities. SDP ID 1 was created with the far-end node 10.20.1.2. Epipe ID 101 was created for customer ID 1 which uses the SDP ID 1.

```
A:ALA-7210M>config>service#
----------------------------------------
...
        sdp 1 mpls create
            description "Default sdp description"
            far-end 10.20.1.2
            lsp "lsp_1_to_B"
            signaling tldp
            no vlan-vc-etype
            path-mtu 9194
            no adv-mtu-override
            keep-alive
                shutdown
                hello-time 10
                hold-down-time 10
                max-drop-count 3
                timeout 5
                no message-length
            exit
            no collect-stats
            no accounting-policy
            no shutdown
        exit
...
    epipe 101 customer 1 vpn 101 create
            description "Default epipe description for service id 101"
            service-mtu 9194
            sap lag-2:101 create
                description "Default sap description for service id 101"
                no tod-suite
                dot1ag
                exit
                ingress
                    qos 1
                    no filter
                exit
            spoke-sdp 101:101 vc-type ether create
                no vlan-vc-tag
                ingress
                    no vc-label
                exit
                egress
                    no vc-label
                exit
                no control-word
                no
                dot1ag
                    mep 1 domain 5 association 101 direction down
                        ccm-enable
                        no ccm-ltm-priority
```

```
                        low-priority-defect remErrXcon
                        no mac-address
                        no shutdown
                    exit
                    mep 1 domain 6 association 101 direction down
                        ccm-enable
                        no ccm-ltm-priority
                        low-priority-defect remErrXcon
                        no mac-address
                        no shutdown
                    exit
                exit
                no collect-stats
                no accounting-policy
                no precedence
                no shutdown
            exit
            no shutdown
...
----------------------------------------
A:ALA-7210M>config>service#
```

Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure a customer account and an SDP. SDP configuration is not needed for 7210 SAS devices configured in Access Uplink mode.

Configuring Customers

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

- Description
- Contact name
- Telephone number

Customer Information

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
                phone phone-number
```

The following displays a basic 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 an SDP

Note: SDPs are not supported by 7210 SAS devices configured in Access Uplink mode.

The most basic SDP must have the following:

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

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 MPLS.
- Each distributed service must have an SDP defined for every remote router to provide VLL, VPLS, and VPRN services.
- A distributed service must be bound to an SDP. By default, no SDP is associated with a service. 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 7210 SAS-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 7210 SAS-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. Only MPLS encapsulation is supported.

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

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

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

Use the following CLI syntax to create an MPLS SDP:

CLI Syntax:

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

The following displays an LSP-signalled MPLS SDP configuration.

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

Configuring a Mixed-LSP SDP

Use the following command to configure an SDP with mixed-LSP mode of operation:

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

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

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

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

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

A special value of the timer dictates that the SDP must never revert to another higher priority LSP type unless the currently active LSP type is down:

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

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

Ethernet Connectivity Fault Management (ETH-CFM)

Ethernet Connectivity Fault Management (ETH-CFM) is defined in two similar standards: IEEE 802.1ag and ITU-T Y.1731. They both specify protocols, procedures, and managed objects to support transport fault management, including discovery and verification of the path, detection and isolation of a connectivity fault for each Ethernet service instance. CFM functionalities are supported on 7210 SAS platforms.

The configuration is split into multiple areas. There is the base ETH-CFM configuration which defines the different Management constructs and administrative elements. This is performed in the ETH-CFM context. The individual management points are configure within the specific service contexts in which they are applied.

The 7210 SAS Services Guide provides the basic service applicable material to build the service specific management points, MEPs and MIPs.

The different service types support a subset of the features from the complete ETH-CFM suite.

ETH-CC used for continuity is available to all MEPs configured within a service. 7210 SAS-M support Down MEPs and UP MEPs, though the support is not available on all platforms. For more information, see the table below.

NOTE: UP MEPs cannot be created by default on system bootup. The user needs to explicitly allocate hardware resources for use with UP MEP feature, using the commands that appear under *configure> system> resource-profile* CLI context. Only after resources have been allocated by the user, UP MEPs are allowed to be created. Until resources are not allocated to UP MEP, the software fails all attempts to create an UP MEP.The troubleshooting tools ETH-LBM/LBR, LTM/LTR ETH-TST defined by the IEEE 802.1ag specification and the ITU-T Y.1731 recommendation are applicable to all MEPs (MIPs where appropriate).

The advanced notification function AIS defined by the ITU-T Y.1731 is supported on Epipe services.

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

For a description of the individual features and functions that are supported, see the OAM and Diagnostics Guide.

Acronym	Callout
1DM	One way Delay Measurement (Y.1731)
AIS	Alarm Indication Signal
CCM	Continuity check message
CFM	Connectivity fault management
DMM	Delay Measurement Message (Y.1731)
DMR	Delay Measurement Reply (Y.1731)
LBM	Loopback message
LBR	Loopback reply
LTM	Linktrace message
LTR	Linktrace reply
ME	Maintenance entity
MA	Maintenance association
MA-ID	Maintenance association identifier
MD	Maintenance domain
MEP	Maintenance association end point
MEP-ID	Maintenance association end point identifier
MHF	MIP half function
MIP	Maintenance domain intermediate point
OpCode	Operational Code
RDI	Remote Defect Indication
TST	Ethernet Test (Y.1731)
SLM	Synthetic Loss Message (Y.1731)
SLR	Synthetic Loss Reply (Y.1731)

ETH-CFM capabilities may be deployed in many different Ethernet service architectures. The Ethernet based SAPs and SDP bindings provide the endpoint on which the management points may be created. The basic functions can be used in different services, VPLS and Epipe . The ETH-CFM functionality is also applicable to broadband access networks. Two models of broadband access are shown below to illustrate how ETH-CFM could be deployed in these cases. (Figure 12 and Figure 13).

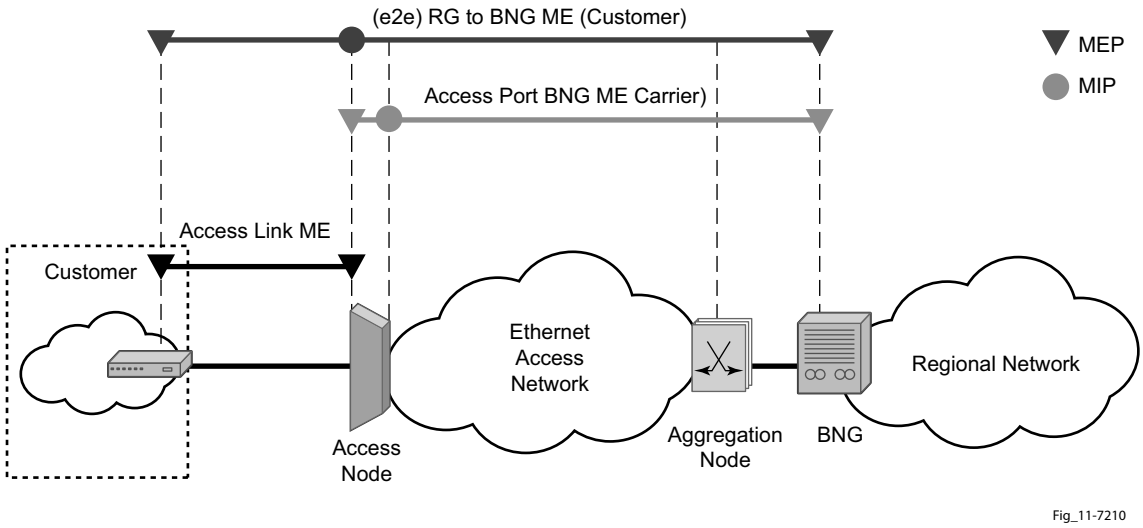

Figure 12: Ethernet OAM Model for Broadband Access - Residential

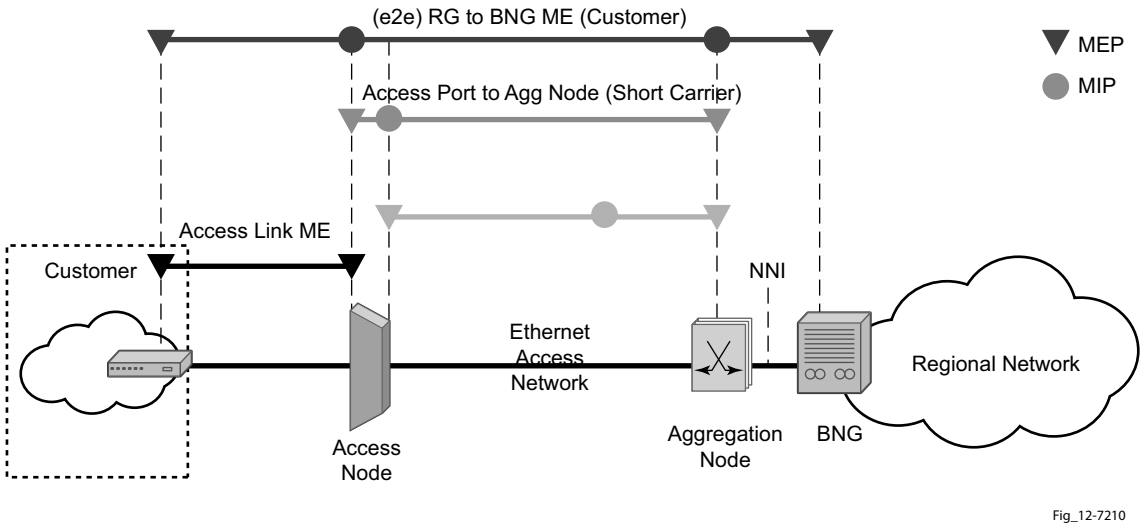

Figure 13: Ethernet OAM Model for Broadband Access - Wholesale

As shown in Figure 17 and Figure 18, the following functions are supported:

- CFM can be enabled or disabled on a SAP or SDP bindings basis.
- The eight ETH-CFM levels are suggested to be broken up numerically between customer 7-5, service provider 4-3 and Operator 2-1. Level 0 is meant to monitor direct connections without any MIPs and should be reserved for port-based facility MEPs. These can be configured, deleted or modified.
- Down MEP and UP MEP with an MEP-ID on a SAP/SDP binding for each MD level can be configured, modified, or deleted. Each MEP is uniquely identified by the MA-ID, MEP-ID tuple.

  → MEP creation on a SAP is allowed only for Ethernet ports (with null, q-tags, qinq encapsulations).
- MIP creation on a SAP for each MD level can be enabled and disabled. MIP creation is automatic or manual when it is enabled. When MIP creation is disabled for an MD level, the existing MIP is removed. For more information on MEP and MIP support, see MEP and MIP Support on page 79

Common Actionable Failures

It is important to note that AIS operates independently from the **low-priority-defect** setting. The **low-priority-defect** setting configuration parameter affects only the ETH-CFM fault propagation and alarming outside the scope of AIS. Any fault in the MEP state machine generates AIS when it is configured. Table 4 illustrates the ETH-CC defect condition groups, configured low-priority-defect setting, priority and defect as it applies to fault propagation.

Table 5: Defect conditions and priority settings

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

MEP and MIP Support

The following is a general table that indicates the ETH-CFM support for the different services and endpoints. It is not meant to indicate the services that are supported or the requirements for those services on the individual platforms.

Table 6: ETH-CFM Support Matrix for 7210 SAS-M Network Mode

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP
Epipe	SAP	Yes	Yes	Ingress MIP and Egress MIP
	SDP	Yes	Yes	Ingress MIP and Egress MIP
VPLS	SAP	Yes	Yes	Ingress MIP only
	Spoke-SDP	Yes	Yes	Ingress MIP only
	Mesh-SDP	Yes	Yes	Not supported
RVPLS	SAP	Not Supported	Not Supported	Not Supported
	IES IPv4 interface	Not Supported	Not Supported	Not Supported
PBB Epipe	I-SAP	Not Supported	Yes	Not Supported
PBB VPLS	I-SAP	Not Supported	Not Supported	Not Supported
PBB B-VPLS	B-SAP	Not Supported	Not Supported	Not Supported
IES	SAP	Not Supported	Not Supported	Not Supported
VPRN	SAP	Not Supported	Not Supported	Not Supported

Table 7: ETH-CFM Support Matrix for 7210 SAS-M Access-Uplink Mode

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP
Epipe	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress MIP and Egress MIP
VPLS	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress MIP and Egress MIP
RVPLS	SAP	Not Supported	Not Supported	Not Supported
	IES IPv4 interface	Not Supported	Not Supported	Not Supported
IES	SAP	Not Supported	Not Supported	Not supported

Table 8: ETH-CFM Support Matrix for 7210 SAS-T Access-Uplink Mode

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP
Epipe	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress MIP and Egress MIP
VPLS	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress MIP only
RVPLS	SAP	Not Supported	Not Supported	Not Supported
	IES IPv4 interface	Not Supported	Not Supported	Not Supported
IES	SAP	Not Supported	Not Supported	Not supported

Table 9: ETH-CFM Support Matrix for 7210 SAS-T Network Mode

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP
Epipe	SAP	Yes	Yes	Ingress MIP and Egress MIP
	SDP	Yes	Yes	Ingress MIP and Egress MIP
VPLS	SAP	Yes	Yes	Ingress MIP only
	Spoke-SDP	Yes	Yes	Ingress MIP only
	Mesh-SDP	Yes	Yes	Not supported
RVPLS	SAP	Not Supported	Not Supported	Not Supported
	IES IPv4 interface	Not Supported	Not Supported	Not Supported
PBB Epipe	I-SAP	Not Supported	Yes	Not Supported
PBB VPLS	I-SAP	Not Supported	Not Supported	Not Supported
PBB B-VPLS	B-SAP	Not Supported	Not Supported	Not Supported
IES	SAP	Not Supported	Not Supported	Not Supported
VPRN	SAP	Not Supported	Not Supported	Not Supported

Note: Ethernet-Rings are not configurable under all service types. Any service restrictions for MEP direction or MIP support will override the generic capability of the Ethernet-Ring MPs. For more information on Ethernet-Rings, refer to the 7210 SAS Interfaces Guide.

Notes:

- On 7210 platforms in some services only ingress MIP functionality is supported. An ingress MIP or a Down MIP processes messages in the ingress direction when the OAM message is received on ingress of the SAP/Port (subject to the other checks). An egress MIP or an UP MIP refers to a MIP that processes message in the egress direction when the OAM message is being sent out of the SAP/port. For more information for service entities support ingress MIP or egress MIP or both, see the tables above.
- On 7210 SAS devices, when two bidirectional MIPs are configured in an Epipe service on both the service entities/endpoints (For example: on both the SAP and the SDP configured in the Epipe service), only the MIP ingressing to the direction of linktrace messages

responds. This is applicable to 7210 SAS platforms that support both ingress and egress MIPs (also referred to as Bi-directional MIPs).

Configuring ETH-CFM Parameters

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

A sample of the global ETH-CFM configuration which defines the domains, associations, linkage o the service id or function, and the globally applicable CCM parameters including the interval and building of the remote MEPs database is shown below.

The following example displays a sample configuration.

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

Defining the MEP and configuring service specific ETH-CFM parameters is performed within the service on the specific SAP or SDP binding. The example using the service VPLS 100 shows this configuration on the SAP.

```
#*A:ALU-7_A>config>service# info
----------------------------------------------
     vpls 100 customer 1 create
          description "VPLS service 100 - Used for MEP configuration example"
               sap 2/2/1:20 create
                    description "2/2/1:20"
                    eth-cfm
                         mep 1 domain 1 association 1 direction down
                              no shutdown
                         exit
                    exit
               exit
          exit
          no shutdown
          exit
          customer 1 create
               description "Default customer"
          exit
          exit
----------------------------------------------
*A:ALU-7_A>config>service#
```

All of the examples shown above were based on IEEE 802.1ag. They are not capable of running Y.1731 functions. To build a Y.1731 context the domain format must be none.

The examples below show the global ETH-CFM configuration and the advanced Y.1731 functions that can be configured. The configuration will reject the configuration of Y.1731 functions within an IEEE 802.1ag context.

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


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

Notes:

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

Applying ETH-CFM Parameters

Apply ETH-CFM parameters to the following entities.

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

CLI Syntax:
```
config>service>epipe>spoke-sdp
 eth-cfm
    mep mep-id domain md-index association ma-index [direction
    {up | down}]
       ccm-enable
       ccm-ltm-priority priority
       eth-test-enable
          test-pattern {all-zeros | all-ones} [crc-enable]
       low-priority-defect {allDef|macRemErrXcon|remErrXcon|
          errXcon|xcon|noXcon}
       [no] shutdown
```

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

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

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

CLI Syntax: oam
eth-cfm linktrace *mac-address* mep *mep-id* domain *md-index* association *ma-index* [ttl *ttl-value*]

eth-cfm loopback *mac-address* mep *mep-id* domain *md-index* association *ma-index* [send-count *send-count*] [size *data-size*] [priority *priority*]

eth-cfm eth-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*] [data-length *data-length*]

eth-cfm one-way-delay-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*]

eth-cfm two-way-delay-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*]

```
eth-cfm two-way-slm-test mac-address mep mep-id domain md-in-
dex association ma-index [priority priority]
```

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] 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

Note : SDPs are not supported by 7210 SAS-M and 7210 SAS-T devices configured in Access Uplink mode.

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

NOTE: 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
```

Layer 2 Control Processing (L2CP)

Operators providing Epipe service need to be able to transparently forward Layer-2 control frames received from the customers. This allows their customers to run these control protocols between the different locations which are part of the L2 VPN service. The 7210 SAS platforms provide user with the following capability:

- An option to tunnel, discard or peer for EFM OAM, LLDP, Dot1x, and LACP.
- BPDU translation and Layer 2 Protocol Tunnelling support for xSTP and CISCO control protocols. This is supported only in a VPLS service. For more information, see the "L2PT and BPDU Translation on page 291".

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

By default, LACP, LLDP, EFM OAM, and Dot1x Layer-2 control protocol untagged packets are discarded if the protocol is not enabled on the port where these frames are received. User has an option to enable peering by enabling the protocol on the port and configuring the appropriate parameters for the protocol. User also has an option to tunnel these packets using an Epipe or VPLS service.

In a VPLS service, the layer-2 control frames are sent out of all the SAPs configured in the VPLS service. It is recommended to use this feature carefully and only when an VPLS is used to emulate an end-to-end Epipe service (that is, an Epipe configured using a 3-point VPLS Service, with one access SAP and 2 access-uplink SAP/SDPs for redundant connectivity). In other words, if the VPLS service is used for multipoint connectivity, it is not recommended to use this feature. When a layer-2 control frame is forwarded out of dot1q SAP or a QinQ SAP, the SAP tags of the egress SAP are added to the packet.

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

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

In addition to the protocols listed above, protocols that are not supported on 7210, For example: GARP, GVRP, ELMI, and others are transparently forwarded in case of a VPLS service. These protocols are transparently forwarded if a NULL SAP, dot1q default SAP , dot1q explicit null SAP or 0.* SAP is configured on the port and received packet is untagged. If the received packet is tagged and matches the tag of any of the SAPs configured on the port, it is forwarded in the

context of the SAP and the service. Else if the received packet is untagged and none of the NULL or dot1q default or dot1q explicit null or 0.* SAP is configured, it is discarded.

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

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

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

Global Services Command Reference

Command Hierarchies

NOTE: All the CLI commands are not available in both access-uplink mode and network modes. Commands applicable to each mode is called out explicitly.

Customer Commands

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

Pseudowire (PW) Commands (applicable only for 7210 SAS devices configured in network mode)

```
config
    — service
        — [no] pw-template  policy-id [use-provisioned-sdp] [create]
            — accounting-policy acct-policy-id
            — no accounting-policy
            — [no] collect-stats
            — [no] control-word
            — [no] disable-learning
            — [no] disable-aging
            — [no] discard-unknown-source
```

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

SDP Commands (Applicable only to 7210 SAS devices configured in network mode)

Note: SDP commands are not applicable for 7210 SAS-M and 7210 SAS-T devices configured in Access Uplink mode.

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

SAP Commands for 7210 SAS devices configured in Network mode

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

SAP Commands for 7210 SAS devices configured in Access-uplink mode

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

ETH-CFM Configuration Commands

config
- **eth-cfm**
  - **domain** *md-index* [**format** *md-name-format*] [**name** *md-name*] **level** *level*
  - **domain** *md-index*
  - **no domain** *md-index*
    - **association** *ma-index* [**format** *ma-name-format*] **name** *ma-name*
    - **association** *ma-index*
    - **no association** *ma-index*
      - [**no**] **bridge-identifier** *bridge-id*
        - **mhf-creation** {**default** | **none** | **explicit**}
        - **no mhf-creation**
        - **vlan** *vlan-id*
        - **no vlan**
      - **ccm-interval** {**100ms** | **1** | **10** | **60** | **600**}
      - **no ccm-interval**
      - [**no**] **remote-mepid** *mep-id*
  - **slm**
    - **[no] inactivity-timer** *timer*

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*]
  - **pw-template [***policy-id***]**
  - **pw-template-using [***policy-id***]**
  - **service-using** [**epipe**][**vpls**][**mirror**][**customer** *customer-id*]
- **eth-ring** [**status**]
- **eth-ring** *ring-index* **hierarchy**
- **eth-ring** *ring-index* [**path** *{a|b}*]
- **eth-cfm**
  - **association** [*ma-index*] [**detail**]
  - **cfm-stack-table** [**port** [*port-id* [**vlan** *vlan-id*]][**level** 0..7] [**direction down**]
  - **cfm-stack-table**
  - **cfm-stack-table port** [{**all-ports**][**level** 0..7][**direction down**]
  - **cfm-stack-table** port-id [**vlan qtag**[.qtag]] [**level** 0..7] **[direction down]**
  - **cfm-stack-table facility [{all-ports|all-lags|all-lag-ports|all-tunnel-meps| all-router-interfaces}] [level** 0..7**] [direction down]**
  - **cfm-stack-table facility lag** *id* **[tunnel** 1..4094**] [level** 0..7**] [direction down]**
  - **cfm-stack-table facility port** *id* **[level 0..7] [direction down]**
  - **cfm-stack-table facility router-interface** *ip-int-name* **[level** 0..7**] [direction down]**
  - **domain** [*md-index*] [**association** *ma-index* | **all-associations**] [**detail**]
  - **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**loopback**] [**linktrace**]
  - **mep** *mep-id* **domain** *md-index* **association** *ma-index* **remote-mepid** *mep-id* | **all-remote-mepids**
  - **mep** *mep-id* **domain** *md-index* **association** *ma-index* **eth-test-results** [**remote-peer** *mac-address*]
  - **mep** *mep-id* **domain** *md-index* **association** *ma-index* **one-way-delay-test** [**remote-peer** *mac-address*]
  - **mep** *mep-id* **domain** *md-index* **association** *ma-index* **two-way-delay-test** [**remote-peer** *mac-address*]
  - **mep** *mep-id* **domain** *md-index* **association** *ma-index* **two-way-slm-test [remote-peer** *macaddress***]**

Global Service Configuration Commands

Generic Commands

shutdown

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

Context config>dot1ag>mep
config>service>sdp
config>service>sdp>keep-alive

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>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* [**create**]
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 *contact-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 none

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.

Pseudowire Commands

pw-template

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

Context config>service

Description This command configures an SDP template.

Parameters *use-provisioned-sdp* — Specifies whether to use an already provisioned SDP. When specified, the tunnel manager will be consulted for an existing active SDP. Otherwise, the default SDP template will be used to use for instantiation of the SDP.

create — This keyword is required when first creating the configuration context. Once the context is created, it is possible to navigate into the context without the create keyword.

control-word

Syntax **[no] control-word**

Context config>service>pw-template

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

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

Default no control-word

SDP Commands

Note: SDP commands are not applicable for 7210 SAS-M and 7210 SAS-T devices configured in Access-Uplink mode.

sdp

Syntax **sdp** *sdp-id* [**mpls**] [**create**]
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 7210 SAS M to a particular service without having to specifically define far end SAPs. Each SDP represents a method to reach a 7210 SAS M router.

The other method is Multi-Protocol Label Switching (MPLS) encapsulation. A 7210 SAS M supports both signaled and non-signaled Label Switched Paths (LSPs) through the network. Non-signaled paths are defined at each hop through the network. Signaled paths are communicated by protocol from end to end using Resource ReserVation Protocol (RSVP). Paths may be manually defined or a constraint-based routing protocol (such as OSPF-TE or CSPF) can be used to determine the best path with specific constraints. An LDP LSP can also be used for an SDP when the encapsulation is MPLS. The use of an LDP LSP type or an RSVP/Static LSP type are mutually exclusive except when the mixed-lsp option is enabled on the SDP.

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

If *sdp-id* does not exist, a new SDP is created. When creating an SDP, the **mpls** keyword must be specified. SDPs are created in the admin down state (**shutdown**) and the **no shutdown** command must be executed 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, the **mpls** keyword is specified. If a keyword is specified for an existing *sdp-id*, an error is generated and the context of the CLI 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

accounting-policy

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

Context config>service>sdp
config>service>pw-template

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
config>service>pw-template

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 no collect-stats

discard-unknown-source

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

Context config>service>pw-template

Description When this command is enabled, packets received with an unknown source MAC address will be dropped only if the maximum number of MAC addresses have been reached. When disabled, the packets are forwarded based on the destination MAC addresses.

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

Default no discard-unknown

limit-mac-move

Syntax **limit-mac-move [blockable | non-blockable]**
no limit-mac-move

Context config>service>pw-template

Description This command indicates whether or not the mac-move agent will limit the MAC re-learn (move) rate.

Default blockable

Parameters *blockable* — The agent will monitor the MAC re-learn rate, and it will block it when the re-learn rate is exceeded.

non-blockable — When specified, a SAP will not be blocked, and another blockable SAP will be blocked instead.

vc-type

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

Context config>service>pw-template

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

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

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

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

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

vlan — Defines the VC type as VLAN. The ethernet and vlan keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings.

vlan-vc-tag

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

Context config>service>pw-template

Description This command specifies an explicit dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured dot1q tag can be overridden by a received TLV specifying the dot1q value expected by the far end. This signaled value must be stored as the remote signaled dot1q value for the binding. The provisioned local dot1q tag must be stored as the administrative dot1q value for the binding.

When the dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.

The no form of this command disables the command

Default no vlan-vc-tag

Parameters *0..4094* — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

adv-mtu-override

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

Context config>service>sdp

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

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

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

The no form of the command disables the VC-type MTU override and returns to the default behavior.

Default no adv-mtu-override

bgp-tunnel

Syntax [**no**] bgp-tunnel

Context config>service>sdp

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

NOTE: The 7210 SAS provides an option to install labels for only those BGP 3107 labelled routes which are in use by services. For more details about this option, see the 7210 SAS-M,T,X,R6 Routing Protocols User Guide.

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

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

far-end

Syntax **far-end** *ip-address* **node-id** *node-id [***global-id** *global-id]*
no far-end

Context config>service>sdp

Description Platforms Supported: MPLS-TP is supported only on 7210 SAS-R6.

This command configures the system IP address of the far-end destination 7210 SAS M7210 SAS X router for the Service Distribution Point (SDP) that is the termination point for a service.

The far-end IP address must be explicitly configured. The destination IP address must be a 7210 SAS M7210 SAS X system IP address.

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

If the SDP uses MPLS encapsulation, the far-end ip-address is used to check LSP names when added to the SDP. If the "to IP address" defined within the LSP configuration does not exactly match the SDP far-end ip-address, the LSP will not be added to the SDP and an error will be generated. Alternatively, and SDP that uses MPLS can have an MPLS-TP node with an MPLS-TP node-id and (optioanlly) global-id. In this case, the SDP must use an MPLS-TP LSP and the SDP signaling parameter must be set to off.

An SDP cannot be administratively enabled until a far-end ip-address or MPLS-TP node-id is defined. The SDP is operational when it is administratively enabled (no shutdown) and the far-end ip-address is contained in the IGP routing table as a host route. OSPF ABRs should not summarize host routes between areas. This can cause SDPs to become operationally down. Static host routes (direct and indirect) can be defined in the local dev ice 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 7210 SAS-M and SAS-X for the SDP in dotted decimal notation.

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

global-id *global-id* — The MPLS-TP Global ID of the far-end system for the SDP, in an unsigned 32-bit integer (0 – 4294967295). This parameter is optonal for an SDP using an MPLS-TP LSP. If note entered, a default value for the Global ID of '0' is used. A global ID of '0' indicates that the far end node is in the same domain as the local node. The user must explicitly configure a Global ID if its value is non-zero.

metric

Syntax **metric** *metric*
no metric

Context config>service>sdp

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

mixed-lsp-mode

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

Context config>service>sdp

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

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

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

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

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

manager.

In the case of the RSVP/LDP SDP, the service manager will program the NHLFE(s) for the active LSP type preferring the RSVP LSP type over the LDP LSP type. If no RSVP LSP is configured or all configured RSVP LSPs go down, the service manager will re-program the linecard with the LDP LSP if available. If not, the SDP goes operationally down.

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

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

Note: For more information on, "***configure>router>ldp>tunnel-down-damp-time*"** command, see 7210 SAS M, X OS MPLS Guide.

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

In the case of the LDP/BGP SDP, the service manager will prefer the LDP LSP type over the BGP LSP type. The service manager will re-program the linecard with the BGP LSP if available otherwise it brings down the SDP operationally.

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

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

Default no mixed-lsp-mode

revert-time

Syntax **revert-time {***revert-time* | *infinite}*
no revert-time

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

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

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

Default 0

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

Values 0 — 600

infinite — This keyword forces the SDP to never revert to another higher priority LSP type unless the currently active LSP type is down.

ldp

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

Context config>service>sdp

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. The above rules are relaxed when the mixed-lsp option is enabled on the SDP.

Default no ldp (disabled)

lsp

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

Context config>service>sdp

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

In MPLS SDP configurations *either* one LSP can be specified.

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

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

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

Default none

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.

signaling

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

Context config>service>sdp

Description This command specifies the signaling protocol used to obtain the ingress and egress pseudowire labels in frames transmitted and received on the SDP. When signaling is *off* 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, MPLS (RSVP or LDP).

tldp — Ingress and egress pseudowire signaling using T-LDP is enabled.

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 device router without packet dropping or IP fragmentation overriding the SDP-type default path-mtu.

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

If the physical **mtu** on an egress interface indicates the next hop on an SDP path cannot support the current **path-mtu**, the operational **path-mtu** on that SDP 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 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 This command configures the number of consecutive SDP keepalive failed request attempts or remote replies that can be missed after which the SDP is operationally downed. If the **max-drop-count** consecutive keepalive request messages cannot be sent or no replies are received, the SDP-ID will be brought operationally down by the keepalive SDP monitoring.
The **no** form of this command reverts the **max-drop-count** *count* value to the default settings.

Default **max-drop-count 3**

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>keep-alive

Description This command configures the SDP monitoring keepalive request message length transmitted.
The **no** form of this command reverts the **message-length** *octets* value to the default setting.

Default 0 — The message length should be equal to the SDP'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

ETH-CFM Configuration Commands

eth-cfm

Syntax **eth-cfm**

Context config

Description This command enables the context to configure 802.1ag CFM parameters.

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* **vlan** *vlan-id]*
no mep *mep-id* **domain** *md-index* **association** *ma-index*

Context config>port>ethernet>
config>lag>
config>router>if>

Description This command provisions the maintenance endpoint (MEP).

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

Parameters *mep-id* — Specifies the maintenance association end point identifier.

Values 1 — 81921

md-index — Specifies the maintenance domain (MD) index value.

Values 1 — 4294967295

ma-index — Specifies the MA index value.

Values 1 — 4294967295

vlan-id — Specific to tunnel facility MEPs which means this option is only applicable to the `lag>eth-cfm>` context. Used to specify the outer vlan id of the tunnel.

Values 1 — 4094

ais-enable

Syntax [**no**] **ais-enable**

Context config>port>ethernet>eth-cfm>mep
config>lag>eth-cfm>mep

Description This command enables the reception of AIS messages.

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

client-meg-level

Syntax **client-meg-level** [[*level* [*level* ...]]
no client-meg-level

Context config>port>ethernet>eth-cfm>mep>ais-enable
config>lag>eth-cfm> mep>ais-enable

Description This command configures the client maintenance entity group (MEG) level(s) to use for AIS message generation. Up to 7 levels can be provisioned with the restriction that the client MEG level must be higher than the local MEG level. Only the lowest client MEG level will be used for facility MEPs.

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

Parameters *level* — Specifies the client MEG level.

Values 1 — 7

Default 1

interval

Syntax **interval** {**1** | **60**}
no interval

Context config>port>ethernet>eth-cfm>mep>ais-enable
config>lag>eth-cfm> mep>ais-enable

Description This command specifies the transmission interval of AIS messages in seconds.

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

Parameters **1** | **60 —** The transmission interval of AIS messages in seconds.

Default 1

priority

Syntax **priority** *priority-value*
no priority

Context config>port>ethernet>eth-cfm>mep>ais-enable
config>lag>eth-cfm> mep>ais-enable

Description This command specifies the priority of the AIS messages generated by the node.

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

Parameters *priority-value* — Specify the priority value of the AIS messages originated by the node.

Values 0 — 7

Default 7

ccm-enable

Syntax [**no**] **ccm-enable**

Context config>port>ethernet>eth-cfm>mep
config>lag>eth-cfm>mep

Description This command enables the generation of CCM messages.

The **no** form of the command disables the generation of CCM messages.

ccm-ltm-priority

Syntax **ccm-ltm-priority** *priority*
no ccm-ltm-priority

Context config>port>ethernet>eth-cfm>mep>
config>lag>eth-cfm>mep>
config>router>if>eth-cfm>mep

Description This command specifies the priority of the CCM and LTM messages transmitted by the MEP. Since CCM does not apply to the Router Facility MEP only the LTM priority is of value under that context.

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

Default *priority* — Specifies the priority value

Values 0 — 7

Default 7

eth-test-enable

Syntax [**no**] **eth-test-enable**

Context config>port>ethernet>eth-cfm>mep>
config>lag>eth-cfm>mep>
config>router>if>eth-cfm>mep

Description For this test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands:

oam eth-cfm eth-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*] [data-length *data-length*]

The **no** form of the command disables eth-test capabilities.

test-pattern

Syntax **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
no test-pattern

Context config>port>ethernet>eth-cfm>mep>eth-test>
config>lag>eth-cfm>mep>eth-test>
config>router>if>eth-cfm>mep>eth-test

Description This command specifies the test pattern of the ETH-TEST frames. This does not have to be configured the same on the sender and the receiver.

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

Parameters **all-zeros —** Specifies to use all zeros in the test pattern.

all-ones — Specifies to use all ones in the test pattern.

crc-enable — Generates a CRC checksum.

Default all-zeros

low-priority-defect

Syntax **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}

Context config>port>ethernet>eth-cfm>mep>eth-test>
config>lag>eth-cfm>mep>eth-test>

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm. This setting is also used to determine the fault state of the MEP which, well enabled to do so, causes a network reaction.

Default macRemErrXcon

Values

allDef	DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
macRemErrXcon	Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
remErrXcon	Only DefRemoteCCM, DefErrorCCM, and DefXconCCM
errXcon	Only DefErrorCCM and DefXconCCM
xcon	Only DefXconCCM; or
noXcon	No defects DefXcon or lower are to be reported

mac-address

Syntax **mac-address** *mac-address*
no mac-address

Context config>port>ethernet>eth-cfm>mep>
config>lag>eth-cfm>mep>
config>router>if>eth-cfm>mep

Description This command specifies the MAC address of the MEP.

The **no** form of the command reverts to the MAC address of the MEP back to the default, that of the port, since this is SAP based.

Parameters *mac-address* — Specifies the MAC address of the MEP.

Values 6-byte unicast mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx) of the MEP. Using the all zeros address is equivalent to the no form of this command.

Default no mac-address

domain

Syntax **domain** *md-index* [**format** *md-name-format*] [**name** *md-name*] **level** *level*
domain *md-index*
no domain *md-index*

Context config>eth-cfm

Description This command configures Connectivity Fault Management domain parameters.

The **no** form of the command removes the MD index parameters from the configuration.

Parameters *md-index* — Specifies the Maintenance Domain (MD) index value.

Values 1 — 4294967295

format {**dns** | **mac** | **none** | **string**} **—** Specifies a value that represents the type (format).

Values
- **dns**: Specifies the DNS name format.
- **mac**: X:X:X:X:X:X-u
X: [0..FF]h
u: [0..65535]d
- **none**: Specifies a Y.1731 domain format and the only format allowed to execute Y.1731 specific functions.
- **string** Specifies an ASCII string.

Default string

name *md-name* **—** Specifies a generic Maintenance Domain (MD) name.

Values 1 — 43 characters

level *level* **—** Specifies the integer identifying the maintenance domain level (MD Level). Higher numbers correspond to higher maintenance domains, those with the greatest physical reach, with the highest values for customers' CFM packets. Lower numbers correspond to lower maintenance domains, those with more limited physical reach, with the lowest values for single bridges or physical links.

Values 0 — 7

association

Syntax **association** *ma-index* [**format** *ma-name-format*] **name** *ma-name*
association *ma-index*
no association *ma-index*

Context config>eth-cfm>domain

Description This command configures the Maintenance Association (MA) for the domain.

Parameters *ma-index* — Specifies the MA index value.

Values 1 — 4294967295

format {**icc-based** | **integer** | **string** | **vid** | **vpn-id**} **—** Specifies a value that represents the type (format).

Values

icc-based:	Only applicable to a Y.1731 context where the domain format is configured as none. Allows for exactly a 13 character name.
integer:	0 — 65535 (integer value 0 means the MA is not attached to a VID.)
string:	raw ascii
vid:	0 — 4095
vpn-id:	RFC-2685, *Virtual Private Networks Identifier* xxx:xxxx, where x is a value between 00 and FF. for example 00164D:AABBCCDD

Default integer

name *ma-name* — Specifies the part of the maintenance association identifier which is unique within the maintenance domain name.

Values 1 — 45 characters

bridge-identifier

Syntax [**no**] **bridge-identifier** *bridge-id*

Context config>eth-cfm>domain>association

Description This command configures the service ID for the domain association. The value must be configured to match the *service-id* of the service where MEPs for this association will be created. Note that there is no verification that the service with a matching *service-id* exists. This is not used for facility MEPs as they are not tied to services.

Parameters *bridge-id* — Specifies the bridge ID for the domain association.

Values 1 — 2147483647

mhf-creation

Syntax **mhf-creation** {**default** | **none** | **explicit**}
no mhf-creation

Context config>eth-cfm>domain>association>bridge-identifier

Description This command determines whether to allow automatic MIP creation for the MA.

NOTE: Please refer to ETH-CFM Support Matrix for 7210 SAS-M Network Mode on page 79ETH-CFM Support Matrix for 7210 SAS-T Access-Uplink Mode on page 80for MEP and MIP support available for different services on different platforms

Default none

Parameters *default* — Specifies that MHFs can be created for this VID only on bridge ports through which this VID can pass without the requirement for a MEP at some lower MA level.

none — Specifies that no MHFs can be created for this VID.

explicit — Specifies that MHFs can be created for this VID only on bridge ports through which this VID can pass, and only if a MEP is created at some lower MA level. There must be at least one lower level MEP provisioned on the same SAP.

vlan

Syntax **vlan** *vlan-id*
no vlan

Context config>eth-cfmg>domain>association>bridge-identifier

Description This command configures the bridge-identifier primary VLAN ID. Note that it is informational only, and no verification is done to ensure MEPs on this association are on the configured VLAN.

Parameters *vlan-id* — Specifies a VLAN ID monitored by MA.

Values 0 — 4094

ccm-interval

Syntax **ccm-interval** {**100ms** | **1** | **10** | **60** | **600**}
no ccm-interval

Context config>eth-cfm>domain>association

Description This command configures the CCM transmission interval for all MEPs in the association.

Table 10: CCM transmission interval for 7210 SAS-M and 7210 SAS-T (Network Mode)

MEP Timer Support	7210 SAS-M Network mode	7210 SAS-T Network mode
Service Down MEP	100ms | 1 | 10 | 60 | 600	1 | 10 | 60 | 600
G8032 Down MEP	100ms | 1 | 10 | 60 | 600	100ms | 1 | 10 | 60 | 600
Service UP MEP	1 | 10 | 60 | 600	1 | 10 | 60 | 600

Table 11: CCM transmission interval for 7210 SAS-M and 7210 SAS-T (Access-Uplink Mode)

MEP Timer Support	7210 SAS-M Access-Uplink mode	7210 SAS-T Access-Uplink mode
Service Down MEP	100ms | 1 | 10 | 60 | 600	100ms | 1 | 10 | 60 | 600

MEP Timer Support	7210 SAS-M Access-Uplink mode	7210 SAS-T Access-Uplink mode
G8032 Down MEP	100ms | 1 | 10 | 60 | 600	100ms | 1 | 10 | 60 | 600
Service UP MEP	1 | 10 | 60 | 600	1 | 10 | 60 | 600

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

Default 10 seconds

Parameters **interval —** Specifies the interval between CCM transmissions to be used by all MEPs in the MA.

remote-mepid

Syntax [**no**] **remote-mepid** *mep-id*

Context config>eth-cfm>domain>association

Description This command configures the remote maintenance association end point (MEP) identifier.

Parameters *mep-id* — Maintenance association end point identifier of a remote MEP whose information from the MEP database is to be returned.

Values 1 — 8191

slm

Syntax **slm**

Context config>eth-cfm

Description This is the container that provides the global configuration parameters for ITU-T Synthetic Loss Measurement (ETH-SL).

inactivity-timer

Syntax **inactivity-timer** *timer*
[no] inactivity-timer

Context config>eth-cfm>slm

Description The time the responder keeps a test active. The time between packets exceed this values within a test the responder marks the previous test as complete. The timer treats any new packets from a peer with the same test-id, source-mac and MEP-ID as a new test responding with the sequence number one.

Default 100 seconds

Parameters *timer* — Specifies the amount of time in seconds.

Values 10 100

VLL Services

In This Chapter

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

Topics in this section include:

Circuit Emulation (Cpipe) Services

Note: CES services are supported only 7210 SAS-M in network mode. It requires the use of T1/E1 CES MDA card in the expansion slot available on the 7210 SAS-M platform. CES services is not supported when operating in access-uplink mode.

Cpipe Service Overview

Cpipe service is the Alcatel-Lucent implementation of TDM pseudowire VLL as defined in the IETF PWE3 working group.

The 7210 SAS-M can support TDM circuit applications that are able to transport delay sensitive TDM traffic over a packet network. For example, in case of business that use legacy T1/E1 interfaces, Cpipe services provide transport services. Cpipe services over MPLS tunnels are supported.

The TDM traffic is transported encapsulated in a TDM VLL over the packet switched network (PSN). The entire T1/E1 frame or part of a frame (n × 64 kb/s) is carried as a TDM VLL over the PSN. At the far end, the transport layer frame structure is regenerated when structured circuit emulation is used, or simply forwarded as part of the payload when unstructured circuit emulation is used.

Cpipe Service Modes

Cpipe services support unstructured circuit emulation mode (SAToP) as per RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP),* and structured circuit emulation mode (CESoPSN) for DS1, E1 and n × 64 kb/s circuits as per RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)*.

Unstructured Mode (SAToP)

Structure-agnostic TDM over Packet (SAToP) is an unstructured circuit emulation mode used for the transport of unstructured TDM or structured TDM (where the structure is ignored).

Note: The word agnostic is used in RFC 4553, but it is not used in the literal sense. The meaning of agnostic in this case is .unaware or independent. Therefore, structure-agnostic is used to mean structure-unaware or structure-independent.

As a structure-unaware or structure-independent service, SAToP service does not align to any framing; the framing mode for the port is set to unframed. For structured TDM, SAToP disregards the bit sequence and TDM structure in order to transport the entire signal over a PSN as a pseudowire.

Structured Mode (CESoPSN)

Structure-aware circuit emulation is used for the transport of structured TDM, taking at least some level of the structure into account. By selecting only the necessary n ?64 kb/s timeslots to transport, bandwidth utilization is reduced or optimized (compared to a full DS1 or E1). Full DS1s or E1s can be transported by selecting all the timeslots in the DS1 or E1 circuit. Framing bits (DS1) or FAS (E1) are terminated at the near end and reproduced at the far end.

When CESoPSN with Channel Associated Signaling (CAS) is selected, the ABCD bits are coded into the T1 or E1 multi-frame packets, transported within the TDM PW, and reconstructed in the T1 or E1 multi-frame at the far end for each timeslot. CAS includes four signaling bits (A, B, C, and D) in the messages sent over a voice trunk. These messages provide information such as the dialed digits and the call state (whether on-hook or off-hook).

The mechanism for E1 CAS is described in ITU-T G.732. When configured for E1 CAS, timeslot 17 carries the signaling information for the timeslots used for voice trunking. Each channel requires four signaling bits, so grouping 16 E1 frames into a multi-frame allows the signaling bits for all 30 channels to be trunked.

As shown in Figure 14, timeslot 1 of all frames within the E1 multi-frame is reserved for alignment, alarm indication, and CRC. For Frame 0, timeslot 17 is reserved for multi-frame alignment bits. For the remaining 15 frames, timeslot 17 contains ABCD bits for two channels.

Note: For E1 CAS, timeslots are numbered 1 to 32 on the 7210 SAS.

For T1 CAS, the signaling bits are transferred using Robbed Bit Signaling (RBS), where the least significant bit in the channel is used periodically to transport these bits instead of voice data.

T1 CAS is supported when ESF or SF framing is configured. ESF framing uses a 24-frame multi-frame and transfers all four signaling bits (ABCD). SF framing uses a 12-frame multi-frame and transfers only the AB bits. The signaling bits are carried in the least significant bit of the following frames:

- A bit in frame 6
- B bit in frame 12
- C bit in frame 18
- D bit in frame 24

Table 12 shows the structure of a T1 ESF multi-frame that uses RBS. The structure of a T1 SF multi-frame is based on 12 frames and only the A and B bits are available.

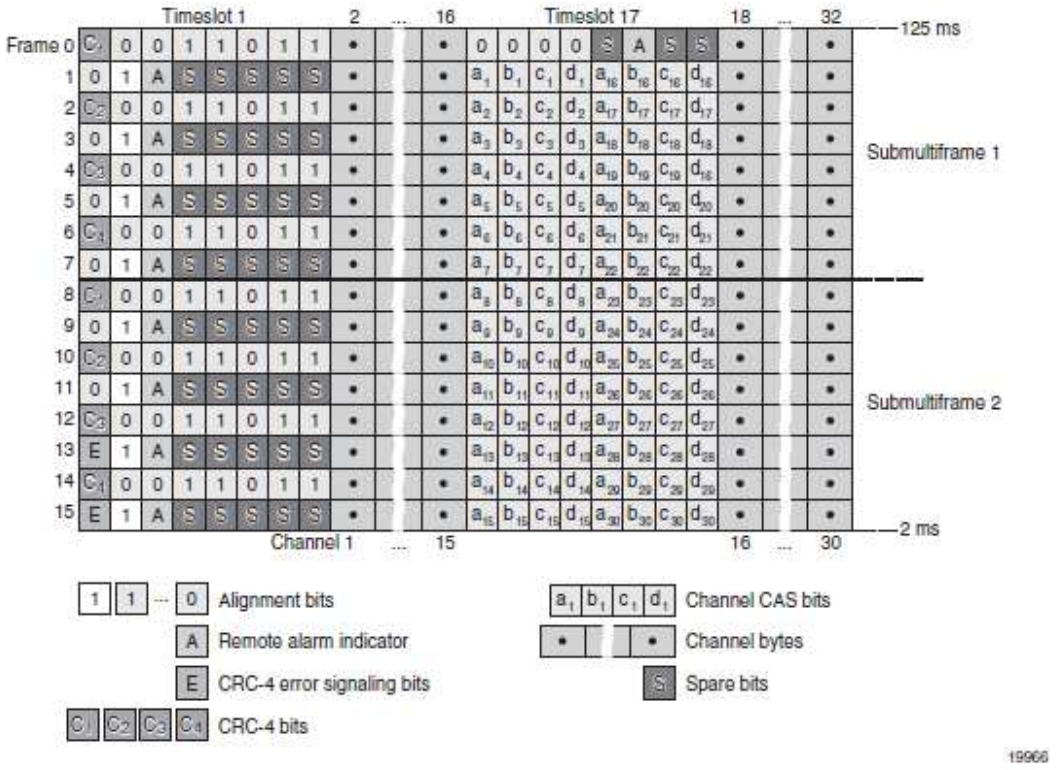

Figure 14: E1 Framing for CAS Support in an E1 Multi-frame

Frame Number	F Bit				Bit Numbers in Each Channel Timeslot		Signaling Channel Designation [4]
	Bit Number within Multiframe	Assignments					
		FAS [1]	DL [2]	CRC [3]	For Character Signal [4]	For Signaling [4]	
1	1	–	m	–	1-8	–	
2	194	–	–	e1	1-8	–	
3	387	–	m	–	1-8	–	
4	580	0	–	–	1-8	–	
5	773	–	m	–	1-8	–	
6	966	–	–	e2	1-7	8	A
7	1159	–	m	–	1-8	–	
8	1352	0	–	–	1-8	–	
9	1545	–	m	–	1-8	–	
10	1738	–	–	e3	1-8	–	
11	1931	–	m	–	1-8	–	
12	2124	1	–	–	1-7	8	B
13	2317	–	m	–	1-8	–	
14	2510	–	–	e4	1-8	–	
15	2703	–	m	–	1-8	–	
16	2896	0	–	–	1-8	–	
17	3089	–	m	–	1-8	–	
18	3282	–	–	e5	1-7	8	C
19	3475	–	m	–	1-8	–	
20	3668	1	–	–	1-8	–	
21	3861	–	m	–	1-8	–	
22	4054	–	–	e6	1-8	–	
23	4247	–	m	–	1-8	–	
24	4440	1	–	–	1-7	8	D

Notes:
1. FAS = frame alignment signal (....001011.....)
2. DL = 4 kb/s data link (m represents message bits)
3. CRC = CRC-6 block check field (e1 to e6 represent check bits)
4. Only applicable for CAS

Table 12: T1 Framing for CAS (RBS) Support in a T1 ESF Multi-frame

TDM Pseudowire Encapsulation

TDM circuits are MPLS-encapsulated as per RFC 4533 (SAToP) and RFC 5086 (CESoPSN), see figures below:

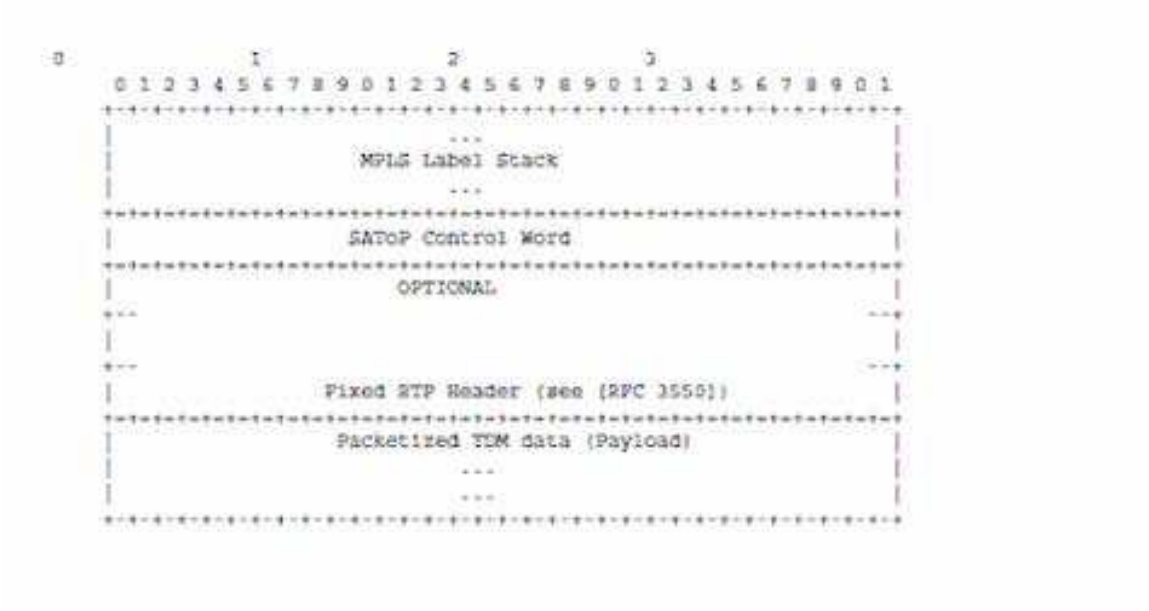

Figure 15: SAToP MPLS Encapsulation

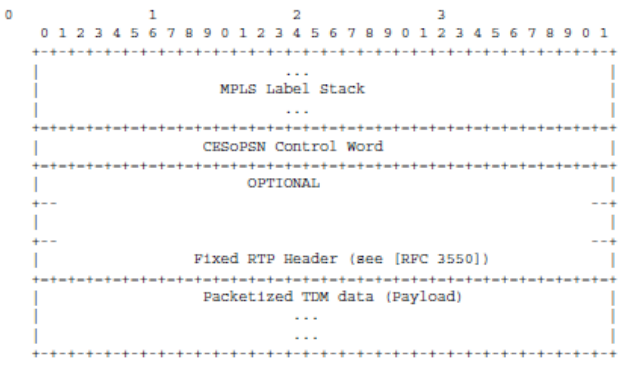

Figure 16: CESoPSN MPLS Encapsulation

Figure 17 shows the format of the CESoPSN TDM payload (with and without CAS) for packets carrying trunk-specific n ??64 kb/s service. In CESoPSN, the payload size is dependent on the number of timeslots used.

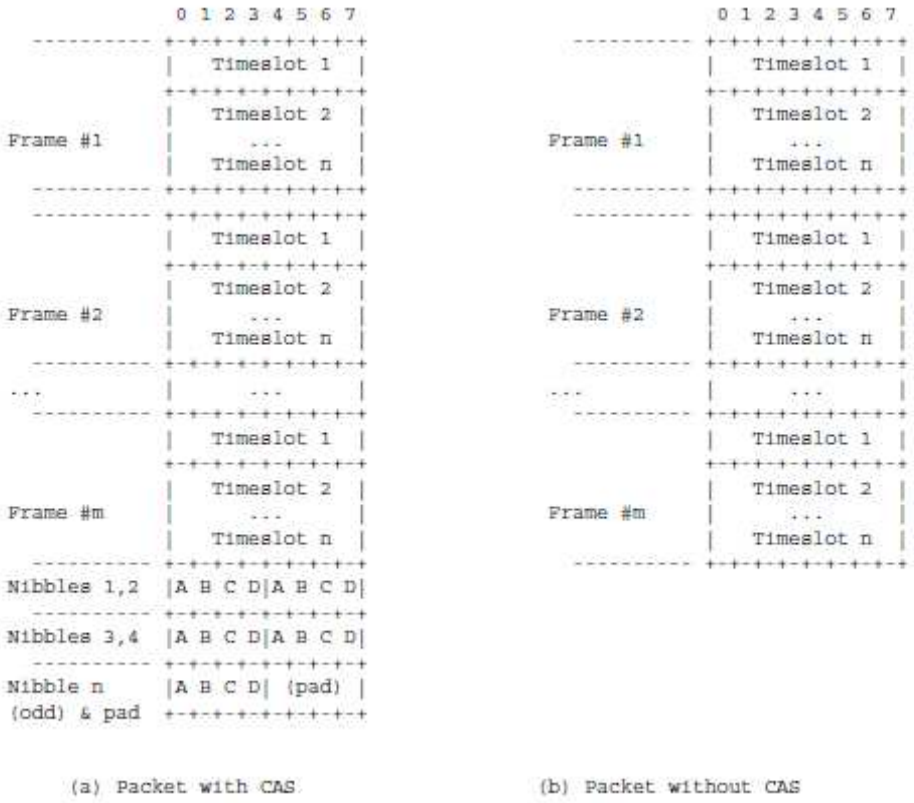

Figure 17: CESoPSN Packet Payload Format for Trunk-Specific n x 64 kb/s (with and without CAS transport)

For CESoPSN without CAS, select the packet size so that an integer number of frames are transported. That is, if n timeslots per frame are to be encapsulated in a TDM PW, then the packet size must be a multiple of n (where n is not equal to 1). For example, if n = 4 timeslots, then the packet size can be 8, 12, 16 and so on.

For CESoPSN with CAS, the packet size is an integer number of frames, where the number of frames is 24 for T1 or 16 for E1, and is not user-configurable. The extra bytes for ABCD (CAS) signaling bits are not included when setting the packet size.

Note: The extra bytes for CAS signaling bits must be included when setting the service-mtu size.

Circuit Emulation Parameters and Options

All ports on the T1/E1 ASAP Adapter card can be configured independently to support TDM circuit emulation across the packet network. Structure-aware mode (CESoPSN) is supported for n × 64 kb/s channel groups in DS1 and E1 circuits. Unstructured mode (SAToP) is supported for full DS1 and E1 circuits. The following parameters and options are described in this section:

- Unstructured
- Structured DS1/E1 CES without CAS
- Structured T1/E1 CES with CAS
- Packet Payload Size
- Jitter Buffer
- RTP Header
- Control Word

Unstructured

Unstructured CES is configured by choosing satop-t1 or satop-e1 as the vc-type when creating a Cpipe service. For DS1 and E1 unstructured circuit emulation, the framing parameter of the port must be set to ds1-unframed and e1-unframed (respectively) because SAToP service ignores the underlying framing. Additionally, channel group 1 must contain all 24 or 32 timeslots, which is configured automatically when channel group 1 is created.

For DS1 and E1 circuit emulation, the payload packet size is configurable and must be an integer value between 64 and 1514 octets and must be a multiple of 32. The payload packet size affects the packet efficiency and packetization delay. Table 13 shows the default values for packet size and packetization delay.

Table 13: Unstructured Payload Defaults

Circuit	Payload Size (Octets)	Packetization Delay (milliseconds)
DS1	192	1.00
E1	256	1.00

Note: When using SAToP to transport DS1 traffic, the framing bit (bit 193) in the DS1 overhead is included and packed in the payload and sent over the PSN. If the underlying framing is ESF, then the Facility Data Link (FDL) channel is transported over the Cpipe as part of the SAToP service. No matter the case, the framing parameter of the port must be set to unframed.

Structured DS1/E1 CES without CAS

Structured CES without CAS is configured by choosing cesopsn as the vc-type when creating a Cpipe service. For n * 64 kb/s structured circuit emulation operation, the framing parameter of the port must be set to a framed setting (such as ESF for DS1). Each channel group contains n DS0s (timeslots), where n is between 1 and 24 timeslots for DS1 and between 1 and 31 timeslots for E1.

The packet payload size is configurable (in octets) and must be an integer multiple of the number of timeslots in the channel group. The minimum payload packet size is 2 octets (based on two frames per packet and one timeslot per frame). See Table 14 for default and minimum payload size values. The maximum payload packet size is 1514 octets.

Each DS1 or E1 frame contributes a number of octets to the packet payload. That number is equal to the number of timeslots configured in the channel group. Thus, a channel group with four timeslots contributes 4 octets to the payload. The timeslots do not need to be contiguous.

Note that a smaller packet size results in a lower packetization delay; however, it increases the packet overhead (when expressed as a percentage of the traffic).

Calculation of Payload Size

The payload size (S), in octets, can be calculated using the following formula:

S = N x F

Where:

N = the number of octets (timeslots) collected per received frame (DS1 or E1)

F = the number of received frames (DS1 or E1) that are accumulated in each CESoPSN packet.

For example, assume the packet collects 16 frames (F) and the channel group contains 4 octets (timeslots) (N). Then the packet payload size (S) is:

S = 4 octets/frame x 16 frames

= 64 octets

Calculation of Packetization Delay

Packetization delay is the time needed to collect the payload for a CESoPSN packet. DS1 and E1 frames arrive at a rate of 8000 frames per second. Therefore, the received frame arrival period is 125 µs.

In the previous example, 16 frames were accumulated in the CESoPSN packet. In this case, the packetization delay (D) can be calculated as follows:

D = 125 µs/frame * 16 frames

= 2.000 ms

Table 14 shows the default and minimum values for frames per packet, payload size, and packetization delay as they apply to the number of timeslots (N) that contribute to the packet payload. The default values are set by the operating system as follows:

- For N = 1, the default is 64 frames/packet
- For 2 ≤?N≤4, the default is 32 frames/packet
- For 5 ≤?N ≤?15, the default is 16 frames/packet
- For N ≥16, the default is 8 frames/packet

	Default Values			Minimum Values		
Number of Timeslots (N)	Frames per Packet (F)	Payload Size (Octets) (S)	Packetization Delay (ms) (D)	Frames per Packet (F)	Payload Size (Octets) (S)	Packetization Delay (ms) (D)
1	64	64	8.000	2	2	0.250
2	32	64	4.000	2	4	0.250
3	32	96	4.000	2	6	0.250
4	32	128	4.000	2	8	0.250
5	16	80	2.000	2	10	0.250
6	16	96	2.000	2	12	0.250
7	16	112	2.000	2	14	0.250
8	16	128	2.000	2	16	0.250
9	16	144	2.000	2	18	0.250
10	16	160	2.000	2	20	0.250
11	16	176	2.000	2	22	0.250

	Default Values			Minimum Values		
Number of Timeslots (N)	Frames per Packet (F)	Payload Size (Octets) (S)	Packetization Delay (ms) (D)	Frames per Packet (F)	Payload Size (Octets) (S)	Packetization Delay (ms) (D)
12	16	192	2.000	2	24	0.250
13	16	208	2.000	2	26	0.250
14	16	224	2.000	2	28	0.250
15	16	240	2.000	2	30	0.250
16	8	128	1.000	2	32	0.250
17	8	136	1.000	2	34	0.250
18	8	144	1.000	2	36	0.250
19	8	152	1.000	2	38	0.250
20	8	160	1.000	2	40	0.250
21	8	168	1.000	2	42	0.250
22	8	176	1.000	2	44	0.250
23	8	184	1.000	2	46	0.250
24	8	192	1.000	2	48	0.250
25	8	200	1.000	2	50	0.250
26	8	208	1.000	2	52	0.250
27	8	216	1.000	2	54	0.250
28	8	224	1.000	2	56	0.250
29	8	232	1.000	2	58	0.250
30	8	240	1.000	2	60	0.250
31	8	248	1.000	2	62	0.250

Table 14: Default and Minimum Payload Size for CESoPSN without CAS

Structured T1/E1 CES with CAS

Structured circuit emulation with CAS is supported for T1 and E1 circuits.

Structured CES with CAS service is configured by choosing cesopsn-cas as the vc-type when creating a Cpipe service. The DS1 or E1 service on the port associated with the Cpipe SAP should be configured to support CAS (via the signal-mode {cas} command) before configuring the Cpipe service to support DS1 or E1 with CAS. Refer to the 7210 SAS Interface Configuration Guide for information on configuring signal mode.

For n *64 kb/s structured circuit emulation with CAS, the implementation is almost identical to that of CES without CAS. When CAS operation is enabled, timeslot 16 cannot be included in the channel group on E1 carriers. The CAS option is enabled or disabled at the port level; therefore, it applies to all channel groups on that E1 port.

The packet size is based on 16 frames per packet for E1 when CAS is enabled and is not user-configurable. For example, if the number of timeslots is 4, then the payload size is 64 octets. This 16-frame fixed configuration is logical because an E1 multi-frame contains 16 frames; therefore, proper bit positioning for the A, B, C, and D CAS signaling bits can be ensured at each end of the pseudo wire. Table 15, “Payload Size for T1 and E1 CESoPSN with CAS,” on page 139 shows the payload sizes based on the number of timeslots.

For CAS, the signaling portion adds (n/2) bytes (n is an even integer) or ((n+1)/2) bytes (n is odd) to the packet, where n is the number of timeslots in the channel group. Note that you do not include the additional signaling bytes in the configuration setting of the TDM payload size. However, the operating system includes the additional bytes in the total packet payload, and the total payload must be accounted for when setting the service-mtu size. Continuing the example above, since n = 4, the total payload is 64 octets plus (4/2 = 2) CAS octets, or 66 octets. Refer to Figure 17 to see the structure of the CES with CAS payload.

CES fragmentation is not supported.

Note: If you configure the service-mtu size to be smaller than the total payload size (payload plus CAS bytes), then the Cpipe will not become operational. This must be considered if you change the service-mtu from its default value.

Table 15: Payload Size for T1 and E1 CESoPSN with CAS

Number of Timeslots	T1			E1		
	Number of Frames per Packet	Payload Size (Octets)	Packetization Delay (ms)	Number of Frames per Packet	Payload Size (Octets)	Packetization Delay (ms)
1	24	24	3.00	16	16	2.00
2	24	48	3.00	16	32	2.00
3	24	72	3.00	16	48	2.00
4	24	96	3.00	16	64	2.00
5	24	120	3.00	16	80	2.00
6	24	144	3.00	16	96	2.00
7	24	168	3.00	16	112	2.00
8	24	192	3.00	16	128	2.00
9	24	216	3.00	16	144	2.00
10	24	240	3.00	16	160	2.00
11	24	264	3.00	16	176	2.00
12	24	288	3.00	16	192	2.00
13	24	312	3.00	16	208	2.00
14	24	336	3.00	16	224	2.00
15	24	360	3.00	16	240	2.00
16	24	384	3.00	16	256	2.00
17	24	408	3.00	16	272	2.00
18	24	432	3.00	16	288	2.00
19	24	456	3.00	16	304	2.00
20	24	480	3.00	16	320	2.00
21	24	504	3.00	16	336	2.00
22	24	528	3.00	16	352	2.00
23	24	552	3.00	16	368	2.00

Number of Timeslots	T1			E1		
	Number of Frames per Packet	Payload Size (Octets)	Packetization Delay (ms)	Number of Frames per Packet	Payload Size (Octets)	Packetization Delay (ms)
24	24	576	3.00	16	384	2.00
25	NA	NA	NA	16	400	2.00
26	NA	NA	NA	16	416	2.00
27	NA	NA	NA	16	432	2.00
28	NA	NA	NA	16	448	2.00
29	NA	NA	NA	16	464	2.00
30	NA	NA	NA	16	480	2.00

Packet Payload Size

The packet payload size defines the number of octets contained in the payload of a TDM pseudowire packet when the packet is transmitted. Each DS0 (timeslot) in a DS1 or E1 frame contributes 1 octet to the payload, and the total number of octets contributed per frame depends on the number of timeslots in the channel group (for example, 10 timeslots contribute 10 octets per frame).

Jitter Buffer

A circuit emulation service uses a jitter buffer to ensure that received packets are tolerant to packet delay variation (PDV). The selection of jitter buffer size must take into account the size of the TDM-encapsulated packets (payload size). A properly configured jitter buffer provides continuous play-out, thereby avoiding discards due to overruns and under runs (packets arriving too early or too late). The maximum receive jitter buffer size is configurable for each SAP configured for circuit emulation. The range of values is from 1 to 250 ms in increments of 1 ms.

Configuration or Design Considerations

Determining the best configuration value for the jitter buffer may require some adjustments to account for the requirements of your network, which can change PDV as nodes are added or removed.

The buffer size must be set to at least three times the packetization delay and no greater than 32 times the packetization delay. Use a buffer size (in ms) that is equal to or greater than the peak-to-

peak packet delay variation (PDV) expected in the network used by circuit emulation service. For example, for a PDV of ±5 ms, configure the jitter buffer to be at least 10 ms.

Note: The jitter buffer setting and payload size (packetization delay) interact such that it may be necessary for the operating system to adjust the jitter buffer setting in order to ensure no loss of packets. Thus, the configured jitter buffer value may not be the value used by the system. Use the **show>service>id service_id>all** command to show the effective PDVT (packet delay variation tolerance).

The following values are the default jitter buffer times for structured circuits, where N is the number of timeslots:

- For N = 1, the default is 32 ms
- For $2 \leq N \leq 4$, the default is 16 ms
- For $5 \leq N \leq 15$, the default is 8 ms
- For $N \geq 16$, the default is 5 ms

Jitter buffer overrun and under run counters are available for statistics and can raise an alarm (optional) while the circuit is operational. For overruns, excess packets are discarded and counted. For under runs, an all-ones pattern is sent for unstructured circuits and an all-ones or a user-defined pattern is sent for structured circuits (based on configuration).

The circuit status and statistics can be displayed using the appropriate show command.

RTP Header

For all circuit emulation channels, the RTP in the header is optional (as per RFC 5086).

When enabled for absolute mode operation, an RTP header is inserted in the MPLS frame upon transmit. Absolute mode is defined in RFC 5086 and means that the ingress PE will set timestamps using the clock recovered from the incoming TDM circuit. When an MPLS frame is received, the RTP header is ignored. The RTP header mode is for TDM pseudowire interoperability purposes only and should be enabled when the other device requires an RTP header.

Control Word

The structure of the control word is mandatory for SAToP and is shown in Figure 18.

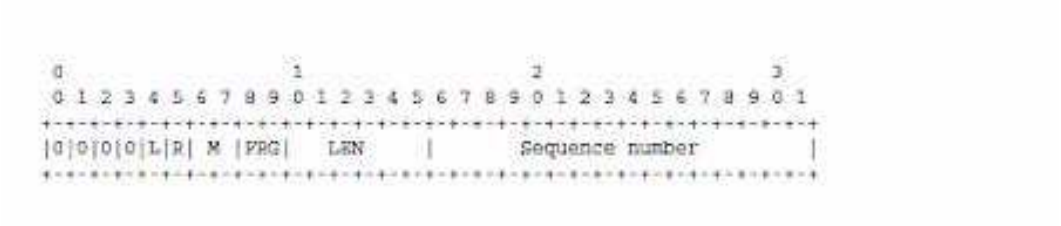

Figure 18: Control Word Bit Structure

The control word descriptions are listed in the Table 16:

Table 16: Control Word Bit Description

Bit(s)	Description
Bits 0 to 3	The use of bits 0 to 3 is described in RFC 4385. These bits are set to ‘0’ unless they are being used to indicate the start of an Associated Channel Header (ACH) for the purposes of VCCV.
L (Local TDM Failure)	The L bit is set to 1 if an abnormal condition of the attachment circuit such as LOS, LOF, or AIS has been detected and the TDM data carried in the payload is invalid. The L bit is cleared (set back to 0) when fault is rectified.
R (Remote Loss of Frames indication)	The R bit is set to 1 if the local CE-bound inter-working function (IWF) is in the packet loss state and cleared (reset to 0) after the local CE-bound IWF is no longer in the packet loss state.
M (Modifier)	The M bits are a 2-bit modifier field. For SAToP, M is set to 00 as per RFC 4553.
Sequence number	The sequence number is used to provide the common pseudowire sequencing function as well as detection of lost packets.

Error Situations

The CE-bound inter-working function (IWF) uses the sequence numbers in the control word to detect lost and incorrectly ordered packets. Incorrectly ordered packets that cannot be re-ordered are discarded.

For unstructured CES, the payload of received packets with the L bit set is replaced with an all-ones pattern. For structured CES, the payload of received packets with the L bit set is replaced with an all-ones or a user-configurable bit pattern. This is configured using the idle-payload-fill command. For structured CES with CAS, the signaling bits are replaced with an all-ones or a user-configurable bit pattern. This is configured using the idle-signal-fill command. Refer to the 7210 SAS Interface Configuration Guide for more information. All circuit emulation services can have a status of up, loss of packets (LOP) or admin down, and any jitter buffer overruns or under runs are logged.

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

Each SAP configuration includes a specific port on which service traffic enters the 7210 SAS router from the customer side (also called the access side). Each port is configured with an encapsulation type. If a port is configured with an IEEE 802.1Q (referred to as Dot1q) encapsulation, then a unique encapsulation value (ID) must be specified.

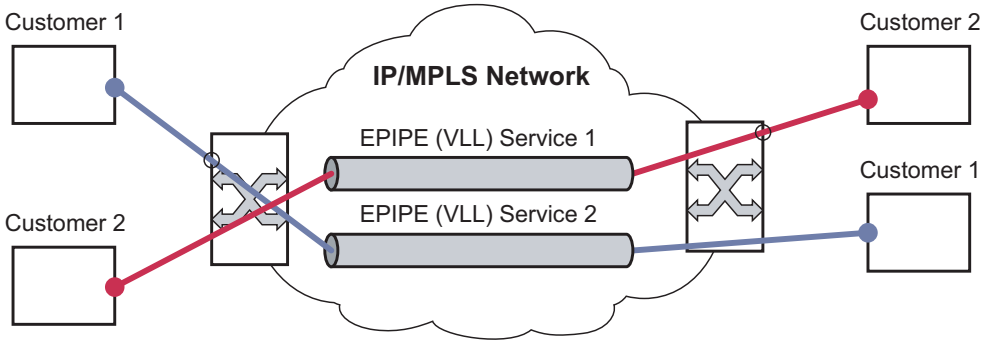

Figure 19: Epipe/VLL Service

Epipe with PBB

A pbb-tunnel may be linked to an Epipe to a B-VPLS. MAC switching and learning is not required for the point-to-point service (all packets ingressing the SAP are PBB encapsulated and forwarded to the PBB tunnel to the backbone destination MAC address and all the packets ingressing the B-VPLS destined for the ISID are PBB de-encapsulated and forwarded to the Epipe SAP. A fully specified backbone destination address must be provisioned for each PBB Epipe instance to be used for each incoming frame on the related I-SAP. If the backbone destination address is not found in the B-VPLS FDB then packets may be flooded through the B-VPLSs

All B-VPLS constructs may be used including B-VPLS resiliency and OAM. Not all generic Epipe commands are applicable when using a PBB tunnel.

Support for processing of packets received with more than 2 tags on a QinQ SAP in Epipe service (only on 7210 SAS devices configured in network mode)

NOTE: 7210 SAS configured in access-uplink mode processes and forwards packets with more than 2 tags. Please see the configuration notes in the Services Chapter for restrictions and use of SAPs in access-uplink mode. This section is applicable only to 7210 SAS devices configured in network mode.

To forward packets with 2 or more tags using a QinQ SAP, a new Epipe service type is available for use when 7210 SAS devices are operating in 'network' mode. This new service will allow for configuration of a QinQ SAP as one endpoint and the following service entities as the other endpoint:

- MPLS spoke-SDP with vc-type set to vc-vlan.
  - → The vc-vlan-tag to be must match the inner-tag VLAN ID value specified in the QinQ SAP.
- dot1q SAP
  - → The VLAN value configured for the dot1q SAP must match the inner-tag VLAN ID value of the QinQ SAP.
- QinQ SAP
  - → The VLAN ID value configured for the innter tag (that is, value of Q1 tag) of the QinQ SAP (that is, Q1.Q2 SAP) must be the same as the inner tag VLAN ID value of the other QinQ SAP.

The device will process the packet as given below in the forward direction:

- If the packet is received on a QinQ SAP, assign an incoming packet to this service based on matching the outermost two tags in the packet header (i.e. in other words the first two tags in the packet header). It will strip only the outermost tag (only a single tag) on ingress and forward the rest on to the other endpoint in the service (see below).
- If the other endpoint the packet is sent out of is a MPLS SDP, then MPLS encapsulation is added.
- If the other endpoint the packet is sent out of is a dot1q SAP packet is forwarded as is, without any egress VLAN checks. It is expected that operator will ensure that the inner tag of the packet matches the dot1q VLAN value.
- If the other endpoint the packet is sent out of is another QinQ SAP (fo example, Q1.Q2 SAP), then another tag (that is, Q2 tag) is added to the packet and sent out of the QinQ SAP.

In the reverse direction, the device will process the packet as given below:

- When traffic is received on the MPLS SDP, the vc-vlan tag is retained as is and the VLAN tag corresponding to the outermost tag configured for the QinQ SAP (i.e. the other endpoint) is added to the packet. The system does not match the vc-vlan tag received in the packet with the configured value (i.e. the inner tag of the QinQ SAP). It is expected that operator will configure both end of the service appropriately to ensure only appropriate packets enter the service.
- When traffic is received on the dot1q SAP, the outermost tag is not stripped and the VLAN tag corresponding to the outermost tag configured for the QinQ SAP is added to the packet.
- If the packet is received on a QinQ SAP, assign an incoming packet to this service based on matching the outermost two tags in the packet header (that is, in other words the first two tags in the packet header). It will strip only the outermost tag (only a single tag) on ingress. The VLAN tag corresponding to the outermost tag configured for the QinQ SAP (that is, the other endpoint) is added to the packet and it is sent out of the QinQ SAP.

Thus, the device processes packets received with 2 or more tags using the MPLS SDP or a dot1q SAP while classifying on the QinQ SAP ingress using 2 tags.

Feature Support, Configuration notes and Restrictions

A new svc-sap-type value "qinq-inner-tag-preserve" is available for configuring the service. This must be used when creating a new Epipe service if this functionality is desired (For example: epipe 10 svc-sap-type qinq-inner-tag-preserve create).

- This service is available only in network mode.
- Epipe service created with the parameter svc-sap-type set to qinq-inner-tag-preserve will allow for only one QinQ SAP and only one SDP of vc-type ‘vc-vlan’. The system will not allow the user to use any other SAP in this new service, that is, NULL SAP, Q1. * SAP, 0.* SAP, etc, are not allowed for configuration in this service. The SDP cannot be of vc-type ‘vc-ether’.
- User can configure vlan-vc-tag value for the SDP, the dot1q SAP VLAN tag value and the inner tag VLAN value of a QinQ SAP to match the VLAN ID value of the inner tag specified in the Q1.Q2 SAP configured in the service (example: if the SAP is 1/1/10:Q1.Q2, then vlan-vc-tag must be set to Q2, the dot1q SAP VLAN value must be Q2, and the inner tag of another QinQ SAP must be set to Q2). If any other value, other than QinQ SAP's inner tag is configured for vlan-vc-tag or dot1q SAP VLAN value, or for the inner tag of the QinQ SAP then it will be errored out by the software. If vlan-vc-tag value is not configured, it defaults to use the inner VLAN tag value. It is highly recommended that the customer configure the vlan-vc-tag value to match the VLAN ID value of the inner tag configured for the QinQ SAP, to avoid misconfiguration.

- Existing QoS and ACL functionality for the Epipe service entities will continue to be available, with the following exceptions:
  - → If the packet is received with more than 2 tags, then IP match-criteria cannot be used with SAP ingress QoS classification and ACLs (both Ingress and Egress ACLs).
  - → If the packet is received with more than 2 tags, then Ethertype value in the mac-criteria cannot be used with SAP ingress QoS classification and ACLs (both Ingress and Egress ACLs).
  - → Dot1p bits from the outermost tag (i.e. Q1 VLAN tag, if the SAP is 1/1/10:Q1.Q2) will be used for SAP ingress classification. Dot1p bits of the outermost tag will be marked on egress, if marking is enabled on the egress port. The Dot1p bit value of the vc-vlan-tag is not used to mark the Dot1p bits of the outermost VLAN tag, when the packets is exiting the QinQ SAP.
- OAM tools
  - → MPLS OAM tools such as vccv-ping, vccv-trace, etc. is supported for the SDPs
  - → Accounting and Statistics for the service entities (e.g. SAP and SDP) will be available as before
  - → CFM/Y.1731 tools are supported. UP and Down MEP is supported on the SAPs and the SDPs configured in the Epipe service.
- Following Redundancy mechanisms available in Epipe service is supported when using MPLS SDP:
  - → Epipe PW redundancy
  - → MC-LAG based protection for access SAPs using the new service type (along with use PW redundancy)

Configuration of Epipe service for processing of packets received with more than 2 tags on a QinQ SAP (only on 7210 SAS devices configured network mode)

The following is the example when the user configures “vlan-vc-tag” value to match the inner tag specified in the Q1.Q2 SAP configured in the service :

```
*A:7210SAS>config>service# info
----------------------------------------------
epipe 10 svc-sap-type qinq-inner-tag-preserve customer 1 create
       sap 1/1/3:10.45 create
    exit
    spoke-sdp 111:69 vc-type vlan create
         vlan-vc-tag 45
    exit
    no shutdown
----------------------------------------------
```

The following is the example of an Epipe service with QinQ SAP and dot1q SAP. In the example below, note that the Dot1q SAP's (1/1/4:45) VLAN value '45', matches the inner tag VLAN value specified with QinQ SAP (1/1/3:10.45).

```
*A:7210>config>service# info
----------------------------------------------
epipe 10 svc-sap-type qinq-inner-tag-preserve customer 1 create
       sap 1/1/3:10.45 create
          no shutdown
    exit
    sap 1/1/4:45 create
         no shutdown
    exit
    no shutdown
exit
----------------------------------------------
```

The following is the example of an Epipe service with 2 QinQ SAPs. In the example below, note that the inner tag of both QinQ SAPs matches and is set to a value of '45'.

```
*A:7210>config>service# info
----------------------------------------------------------------------------------------------
-
epipe 10 svc-sap-type qinq-inner-tag-preserve customer 1 create
     sap 1/1/3:10.45 create
          no shutdown
     exit
     sap 1/1/4:200.45 create
          no shutdown
     exit
     no shutdown
exit
----------------------------------------------------------------------------------------------
-----
```

Epipe Oper State decoupling

An epipe service transitions to an operation state, ‘Down’ when only a single entity SAP or Binding is active and the operation state of the mate is down or displays an equivalent state. The default behavior does not allow operators to validate the connectivity and measure performance metrics. With this feature an option is provided to allow operators to validate the connectivity and measure performance metrics of an epipe service prior to the customer handoff. The operator can also maintain performance and continuity measurement across their network regardless of the connectivity between the terminating node and the customer. If the SAP between the operator and the customer enters a Oper Down state, the epipe remains Operationally UP, so the results can continue to be collected uninterrupted. The operator receives applicable port or SAP alerts/alarms. This option is available only for the customer facing SAP failures. If a network facing SAP or Spoke-SDP fails the operational state of the epipe service is set to 'Down'. In other words, there is no option to hold the service in an UP state, if a network component fails.

The following functionality is supported:

- Configuration under SAP is required to change the default behavior of the epipe service in response to the SAP failure.
- The user can create a SAP on a LAG where the LAG has no port members. In this case, the operator configures the “*ignore-oper-state*” on the SAP and the service remains operational. However, as there are no ports existing in the LAG member group, there is no extraction function that can be created. This feature protects against an established working configuration with full forwarding capabilities from failing to collect PM data. The user should shutdown their equipment and place the epipe SAP in an operationally down state.
- The SAP connecting the provider equipment to the customer is configured to hold the epipe service status UP when the customer facing SAP enters any failed state. Only one SAP per epipe is allowed to be configured.
- Any failure of the network entity (network SAP or SDP-Binding) still cause the epipe service to transition to OPER=DOWN.
- As the service remains operationally up, all bindings should remain operationally up and should be able to receive and transmit data. The PW status represents the failed SAP in the LDP status message, but this does not prevent the data from using the PW as a transport, in or out. This is the same as LDP status messaging.
- The SAP failure continues to trigger normal reactions, except the operational state of the service
- ETH-CFM PM measurement tools (DMM/SLM) can be used with the UP MEP on the failed SAP to collect performance metric. Additionally, CFM troubleshooting tools & connectivity (LBM, LTM, AIS, CCM) can be used and will function normally.

- ETH-CFM CCM processing and fault propagation does not change. Even when a SAP fails with the hold service UP configuration, CCM sets the Interface Status TLV to "*Down*".
- VPLS services remain operationally UP until the final entity in the service enters a failed operational state. There are no changes to VPLS services and the change is specific to epipe.

Pseudowire Switching

Note: 7210 SAS-X and 7210 SAS-R6 device can be configured as S-PE nodes. 7210 SAS-M and 7210 SAS-T nodes can be configured only as T-PE nodes.

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

Services with one SAP and one spoke SDP are created normally on the PE; however, the target destination of the SDP is the pseudowire switching node instead of what is normally the remote PE.

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

In the Figure 20, the user configures a regular Epipe VLL service PE1 and 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 pseudowire switching node. In addition, the user configures an Epipe VLL service on the pseudowire switching node using the two SDPs.

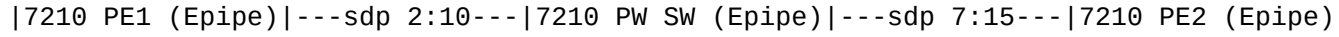

Figure 20: Pseudowire Service Switching Node

Pseudowire Switching with Protection

Pseudowire switching scales VLL and VPLS services over a multi-area network by removing the need for a full mesh of targeted LDP sessions between PE nodes. Figure 21 illustrates the use of pseudowire redundancy to provide a scalable and resilient VLL service across multiple IGP areas in a provider network.

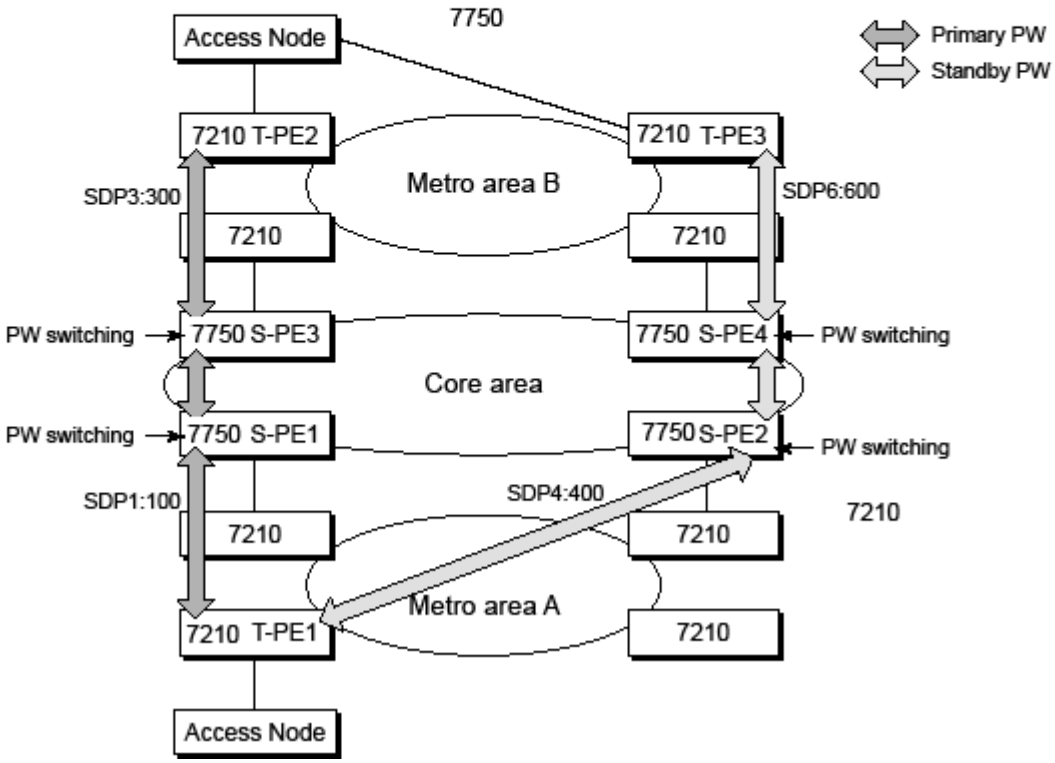

Figure 21: VLL Resilience with Pseudowire Redundancy and Switching

In the network in Figure 21, PE nodes act as masters and pseudowire switching nodes act as slaves for the purpose of pseudowire signaling. A switching node will need to pass the SAP Interface Parameters of each PE to the other.T-PE1 sends a label mapping message for the Layer 2 FEC to the peer pseudowire switching node" for example, S-PE1. It will include the SAP interface parameters, such as MTU, in the label mapping message. S-PE1 checks the FEC against the local information and if a match exists, it appends the optional pseudowire switching point TLV to the FEC TLV in which it records its system address. T-PE1 then relays the label mapping message to S-PE2. S-PE2 performs similar operations and forwards a label mapping message to T-PE2. The same procedures are followed for the label mapping message in the reverse direction, for example,

from T-PE2 to T-PE1. S-PE1 and S-PE2 will effect the spoke SDP cross-connect only when both directions of the pseudowire have been signaled and matched.

The pseudowire switching TLV is useful in a few situations. First, it allows for troubleshooting of the path of the pseudowire especially if multiple pseudowire switching points exist between the two T-PE nodes. Secondly, it helps in loop detection of the T-LDP signaling messages where a switching point receives back a label mapping message it already relayed. Finally, it can be inserted in pseudowire status messages when they are sent from a pseudowire switching node towards a destination PE.

Pseudowire status messages can be generated by the T-PE nodes. Pseudowire status messages received by a switching node are processed and then passed on to the next hop. An S-PE node appends the optional pseudowire switching TLV, with its system address added to it, to the FEC in the pseudowire status notification message only if it originated the message or the message was received with the TLV in it. Otherwise, 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 VCID value in the FEC TLV.

Pseudowire Switching Behavior

In the network in Figure 21, 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.

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

- When the local status for both spokes is up, the S-PE passes any received SAP or SDP-binding generated status notification message unchanged, for example, the status notification TLV is unchanged but the VC-ID in the FEC TLV is set to 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:

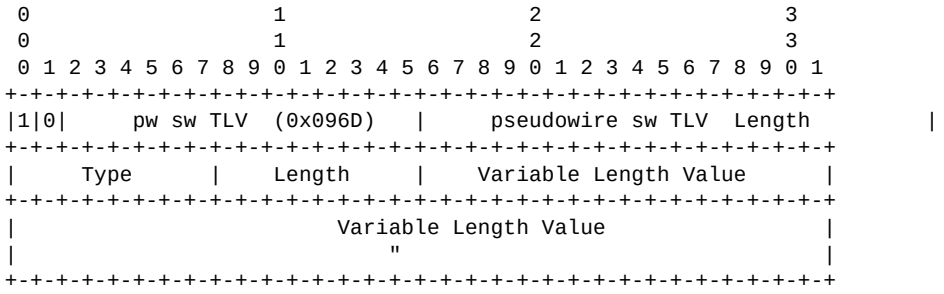

- PW sw TLV Length — Specifies the total length of all the following pseudowire switching point TLV fields in octets
- Type — Encodes how the Value field is to be interpreted.
- Length — Specifies the length of the Value field in octets.
- Value — Octet string of Length octets that encodes information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

Static-to-Dynamic Pseudowire Switching

When one segment of the pseudowire cross-connect at the S-PE is static while the other is signaled using T-LDP, the S-PE operates much like a T-PE from a signaling perspective and as an S-PE from a data plane perspective.

The S-PE signals a label mapping message as soon as the local configuration is complete. The control word C-bit field in the pseudowire FEC is set to the value configured on the static spokesdp.

When the label mapping for the egress direction is also received from the T-LDP peer, and the information in the FEC matches that of the local configuration, the static-to-dynamic crossconnect is effected.

Note that it is possible that end nodes of a static pseudowire segment be misconfigured. In this case, an S-PE or T-PE node may be receiving packets with the wrong encapsulation. In this case, it is possible that an invalid payload will be forwarded over the pseudowire or the SAP respectively. Furthermore, if the S-PE or T-PE node is expecting the control word in the packet encapsulation and the received packet comes with no control word but the first nibble below the label stack is 0x0001, the packet may be mistaken for a VCCV OAM packet and may be forwarded to the CPM. In that case, the CPM will perform a check of the IP header fields such as version, IP header length, and checksum. If any of this fails the VCCV packet will be discarded.

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

Pseudowire and VPLS link redundancy extends link-level resiliency for pseudowires and VPLS to protect critical network paths against physical link or node failures. These innovations enable the virtualization of redundant paths across the metro or core IP network to provide seamless and transparent fail-over for point-to-point and multi-point connections and services. When deployed with multi-chassis LAG, the path for return traffic is maintained through the pseudowire or VPLS switchover, which enables carriers to deliver “always on” services across their IP/MPLS networks.

VLL Resilience with Two Destination PE Nodes

Figure 22 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.

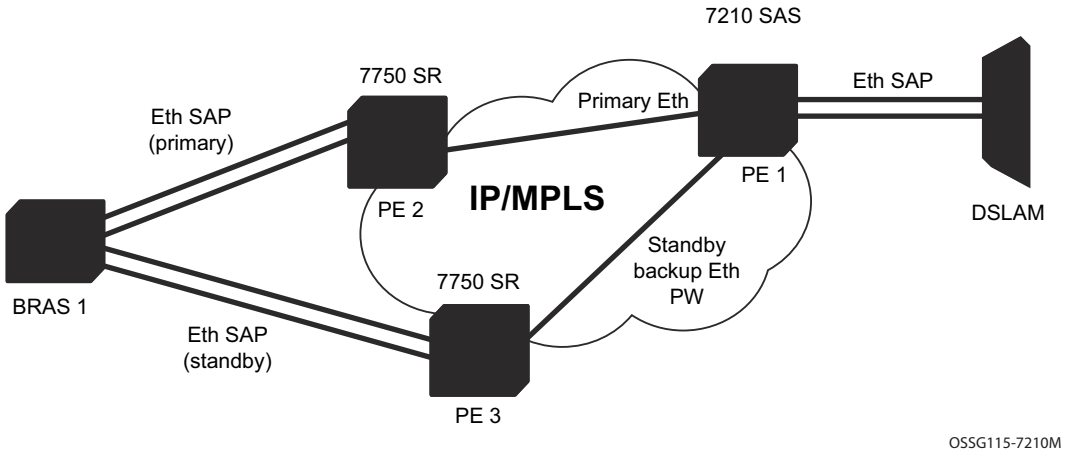

Figure 22: VLL Resilience

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.

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.

Alcatel-Lucent's routers support the ability to configure multiple secondary standby pseudowire paths. For example, PE1 uses the value of the user configurable precedence parameter associated with each spoke SDP to select the next available pseudowire path after the failure of the current active pseudowire (whether it is the primary or one of the secondary pseudowires). 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.

Alcatel-Lucent's routers support the ability for a user-initiated manual switchover of the VLL path to the primary or any of the secondary be supported to divert user traffic in case of a planned outage such as in node upgrade procedures.

Master-Slave Operation

NOTE: 7210 SAS devices support only standby-signaling-master option. 7210 does not support the CLI command standby-signaling-slave. In the discussion below, reference to standby-signaling-slave command is only used to describe the solution. 7210 device can be used only where standby-signaling-master is used in the example below.

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

Figure 23 illustrates the operation of master-slave pseudowire redundancy. In this scenario, an Epipe service is provided between CE1 and CE2. CE2 is dual homed to PE2 and PE3, and thus PE1 is dual-homed to PE2 and PE3 using Epipe spoke SDPs. The objectives of this feature is to ensure that only one pseudowire is used for forwarding in both directions by PE1, PE2 and PE3 in the absence of a native dual homing protocol between CE2 and PE2/PE3, such as MC-LAG. In normal operating conditions (the SAPs on PE2 and PE3 towards CE2 are both up and there are no defects on the ACs to CE2), PE2 and PE3 cannot choose which spoke SDP to forward on based on the status of the AC redundancy protocol.

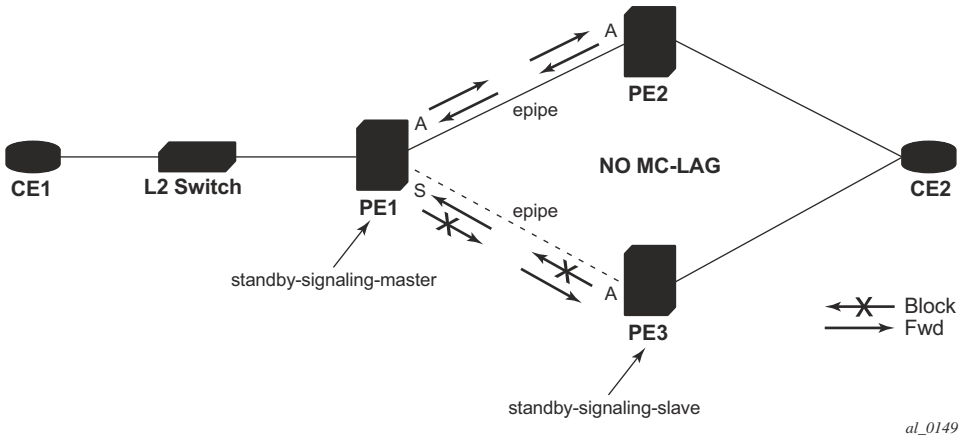

Figure 23: Master-Slave Pseudowire Redundancy

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

This is achieved as follows. The **standby-signaling-master** state is activated on the VLL endpoint in PE1. In this case, a spoke SDP is blocked in the transmit direction at this master endpoint if it is either in operDown state, or it has lower precedence than the highest precedence spoke SDP, or the given peer PE signals one of the following pseudowire status bits:

- Pseudowire not forwarding (0x01)
- SAP (ingress) receive fault (0x02)
- SAP (egress) transmit fault (0x04)
- SDP binding (ingress) receive fault (0x08)
- SDP binding (egress) transmit fault (0x10)

The fact that the given spoke SDP has been blocked will be signaled to LDP peer through the pseudowire status bit (PW FWD Standby (0x20)). This will prevent traffic being sent over this spoke SDP by the remote peer, but obviously only in case that remote peer supports and reacts to pseudowire status notification. Previously, this applied only if the spoke SDP terminates on an IES, VPRN or VPLS. However, if standby-signaling-slave is enabled at the remote VLL endpoint then the Tx direction of the spoke SDP will also be blocked, according to the rules in Operation of Master-Slave Pseudowire Redundancy with Existing Scenarios on page 164.

Note that although master-slave operation provides bidirectional blocking of a standby spoke SDP during steady-state conditions, it is possible that the Tx directions of more than one slave endpoint can be active for transient periods during a fail-over operation. This is due to slave endpoints

transitioning a spoke SDP from standby to active receiving and/or processing a pseudowire preferential forwarding status message before those transitioning a spoke SDP to standby. This transient condition is most likely when a forced switch-over is performed, or the relative preferences of the spoke SDPs is changed, or the active spoke SDP is shutdown at the master endpoint. During this period, loops of unknown traffic may be observed. Fail-overs due to common network faults that can occur during normal operation, a failure of connectivity on the path of the spoke SDP or the SAP, would not result in such loops in the data path.

Local Rules at Slave VLL PE

It must not be possible to configure standby-signaling-slave on endpoints or spoke SDPs bound to an IES, VPRN, ICB, MC-EP or that form part of an MC-LAG or MC-APS.

If ‘standby-signaling-slave’ is configured on a given spoke SDP or explicit endpoint, then the following rules apply. Note that the rules describe the case of several spoke SDPs in an explicit endpoint. The same rules apply to the case of a single spoke SDP outside of an endpoint where no endpoint exists:

Rules for processing endpoint SAP active/standby status bits:

→ Since the SAP in endpoint X is never a part of a MC-LAG/MC-APS instance, a forwarding status of ACTIVE is always advertised.

Rules for processing and merging local and received endpoint object status Up/Down operational status:

1. 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.
2. 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 “PW not forwarding”, the node must send status bits of “SAP Down” over all ‘Y’ endpoint spoke SDPs.
3. 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.
4. 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.
5. If a spoke SDP in endpoint ‘Y’, received T-LDP “SAP down” status bits, and/or received T-LDP “SDP-binding down” status bits, and/or received status bits of “PW not forwarding”, the node saves this status and takes no further action. The saved status is used for selecting the active transmit endpoint object as per the pseudo-code in Section 5.1.2.
6. If, all objects in endpoint ‘Y’, or a single spoke SDP that exists outside of an endpoint (and no endpoint exists), 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 “PW not forwarding”, and/or the received status bits of ‘PW FWD standby’, the node

must send a “SAP down” notification on the ‘X’ endpoint SAP via the SAP specific OAM signal, if applicable.

7. If the peer PE for a given object in endpoint ‘Y’ signals ‘PW FWD standby’, the spoke SDP must be blocked in the transmit direction and the spoke SDP is not eligible for selection by the active transmit selection rules.
8. If the peer PE for a given object in endpoint ‘Y’ does not signal ‘PW FWD standby’, then spoke SDP is eligible for selection.

Operation of Master-Slave Pseudowire Redundancy with Existing Scenarios

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

VLL Resilience

Figure 24 displays a VLL resilience path example. An sample configuration follows.

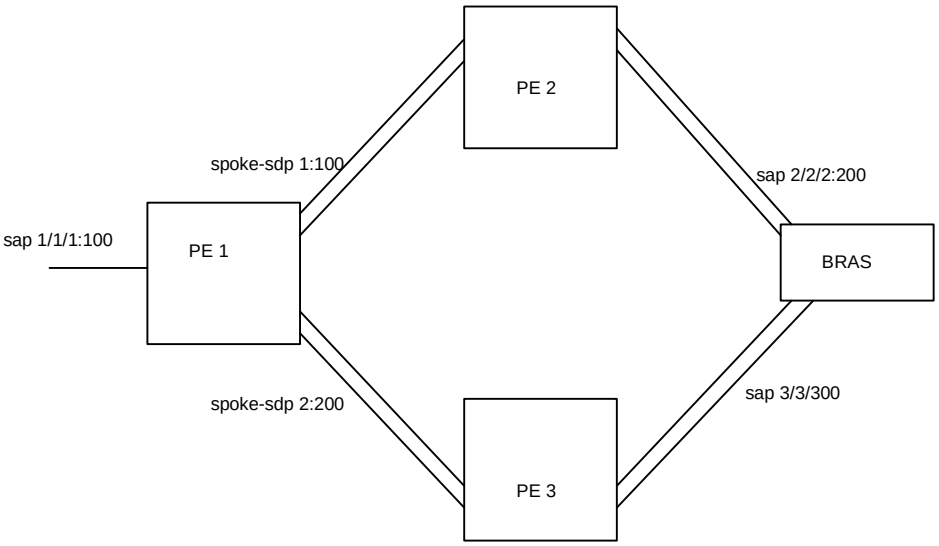

Figure 24: VLL Resilience

Note that a **revert-time** value of zero (default) means that the VLL path will be switched back to the primary immediately after it comes back up

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

```
     spoke-sdp 1:100
          standby-signaling-slave
```

PE3

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

VLL Resilience for a Switched PW Path

Note: 7210 SAS-R6 and 7210 SAS-X nodes can be configured as S-PE nodes. 7210 SAS-M and 7210 SAS-T nodes can act only as T-PE nodes.

Figure 25 displays a VLL resilience for a switched pseudowire path example. A sample configuration follows.

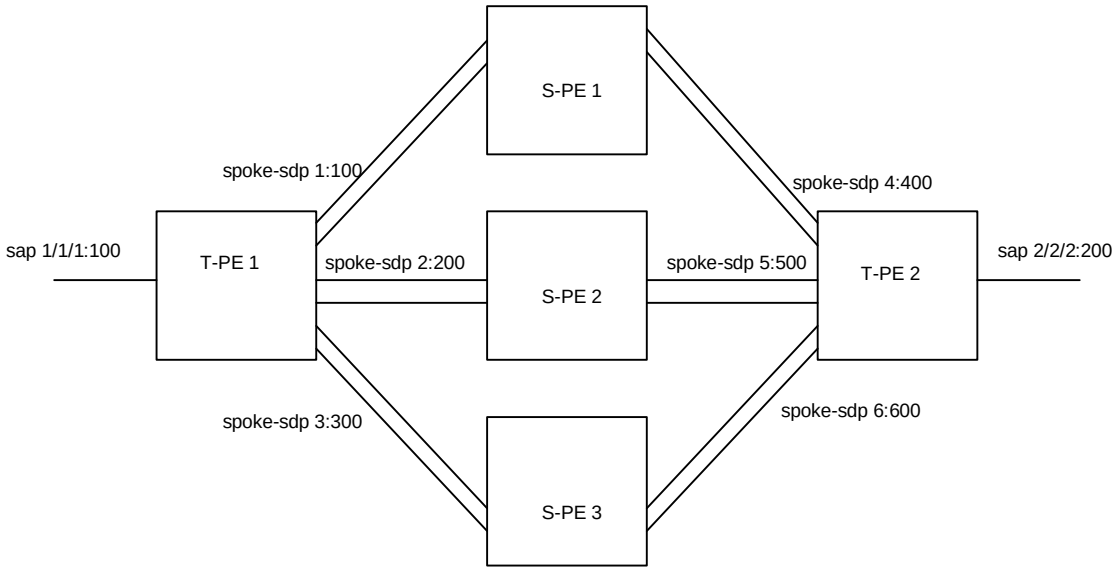

Figure 25: VLL Resilience with Pseudowire Switching

Configuration

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

```
         precedence 1

T-PE2
configure service epipe 1
     endpoint X
     exit
     endpoint Y
     revert-time 100
     standby-signaling-slave
     exit
     sap 2/2/2:200 endpoint X
     spoke-sdp 4:400 endpoint Y
          precedence primary
     spoke-sdp 5:500 endpoint Y
          precedence 1
     spoke-sdp 6:600 endpoint Y
          precedence 1
```

S-PE1

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

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

Access Node Resilience Using MC-LAG and Pseudowire Redundancy

Figure 26 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.

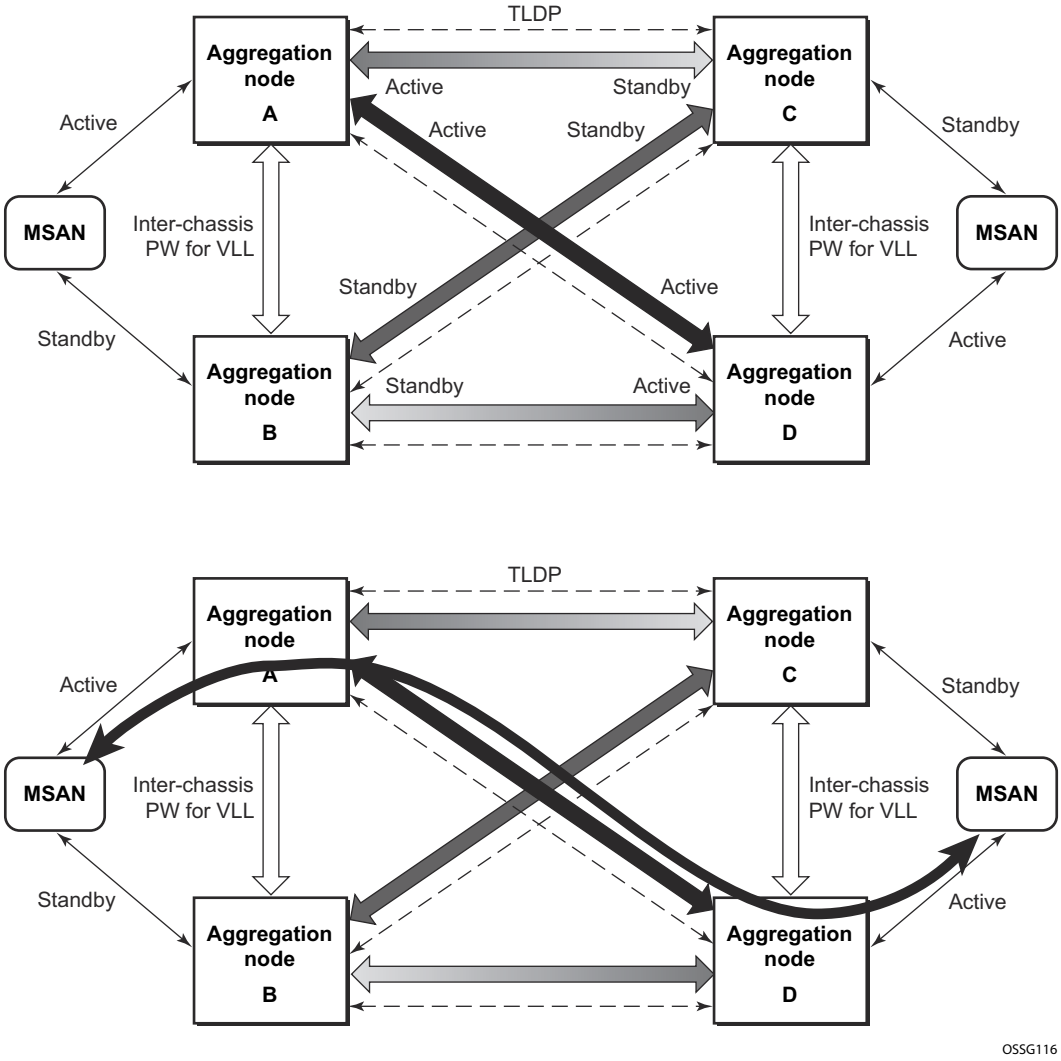

Figure 26: Access Node Resilience

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 7210 SAS aggregation node. All spoke SDPs are of secondary type and there is no use of a primary pseudowire type in this mode of operation. Node A is in the active state according to its local MC-LAG instance and 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.

7210 SAS node selects a pseudowire as the active path for forwarding packets when both the local pseudowire status and the received remote pseudowire status indicate active status. However, a

7210 SAS device in standby status according to the SAP in its local MC-LAG instance is capable of processing packets for a VLL service received over any of the pseudowires which are up. This is to avoid black holing of user traffic during transitions. The 7210 SAS standby node forwards these packets to the active node bye 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 described in the application. A 7210 SAS node uses this to select which pseudowire to forward packet to in case both pseudowires show active/active for the local/remote status during transitions.

Only VLL service of type Epipe is supported in this application. 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.

VLL Resilience for a Switched Pseudowire Path

Figure 27 illustrates the use of both pseudowire redundancy and pseudowire switching to provide a resilient VLL service across multiple IGP areas in a provider network.

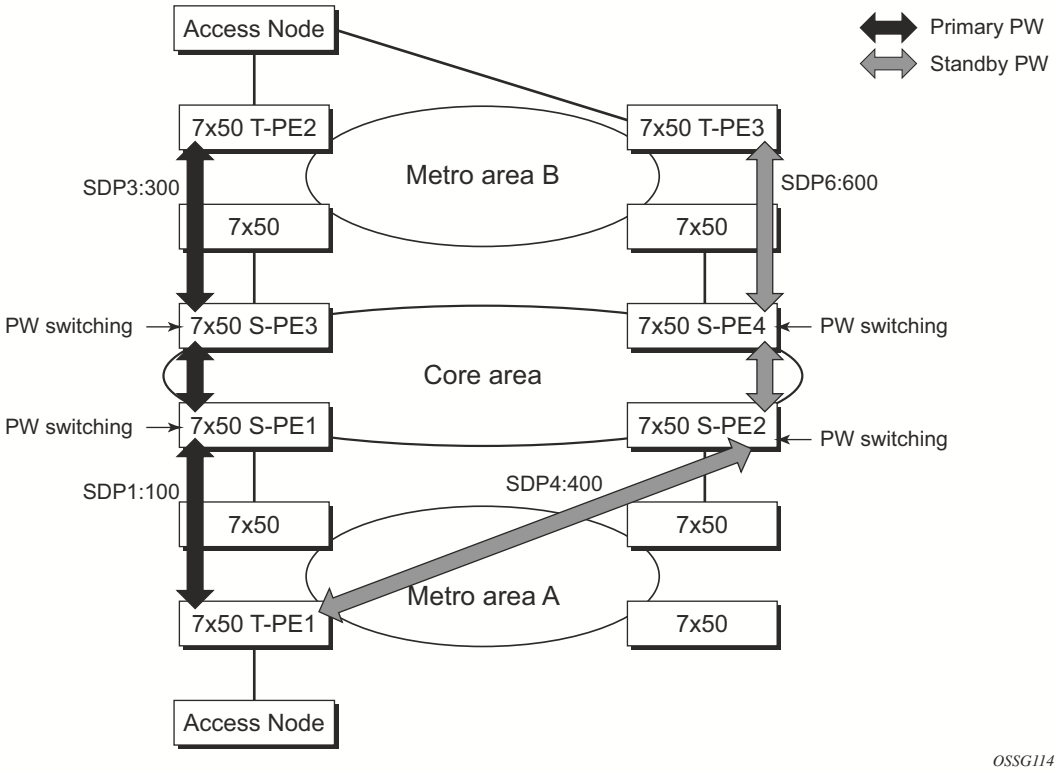

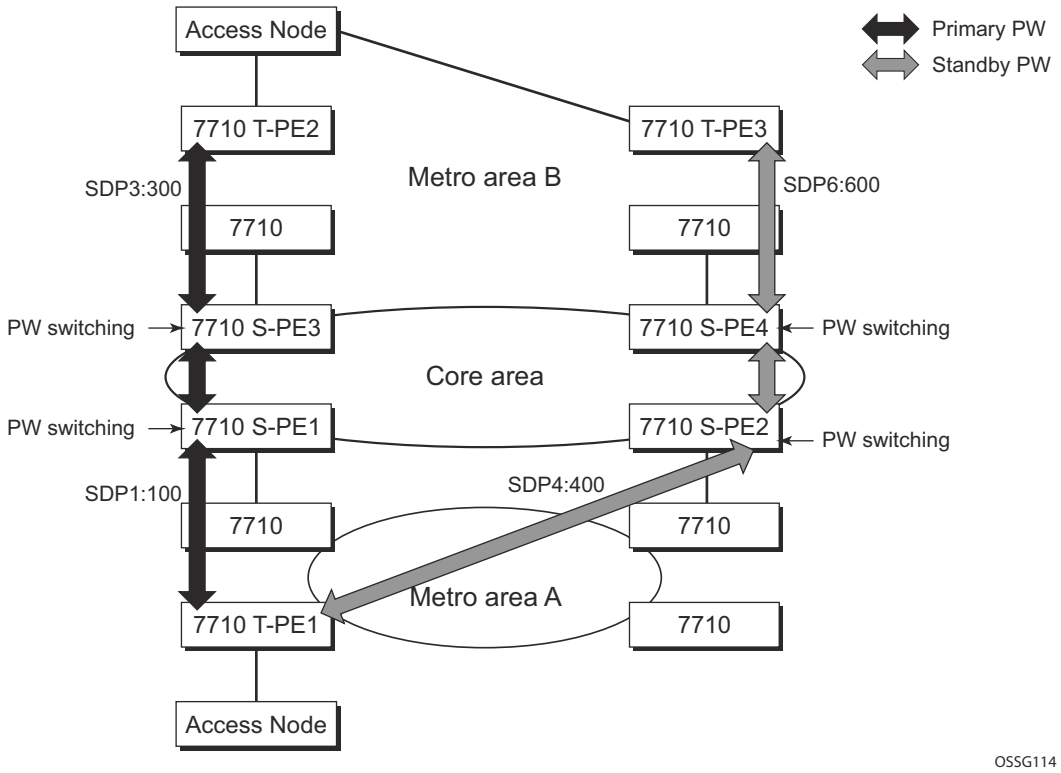

Figure 27: VLL Resilience with Pseudowire Redundancy and Switching

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 PE nodes as the number of these nodes grows over time.

Like in the application in VLL Resilience with Two Destination PE Nodes on page 339, the T-PE1 node switches the path of a VLL to a secondary standby pseudowire in the case of a network side failure causing the VLL binding status to be DOWN or if T-PE2 notified it that the remote SAP went down. This application requires that pseudowire status notification messages generated by either a T-PE node or a S-PE node be processed and relayed by the S-PE nodes.

Note 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.

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.

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 28 illustrates the model for a redundant VLL service based on the concept of endpoints.

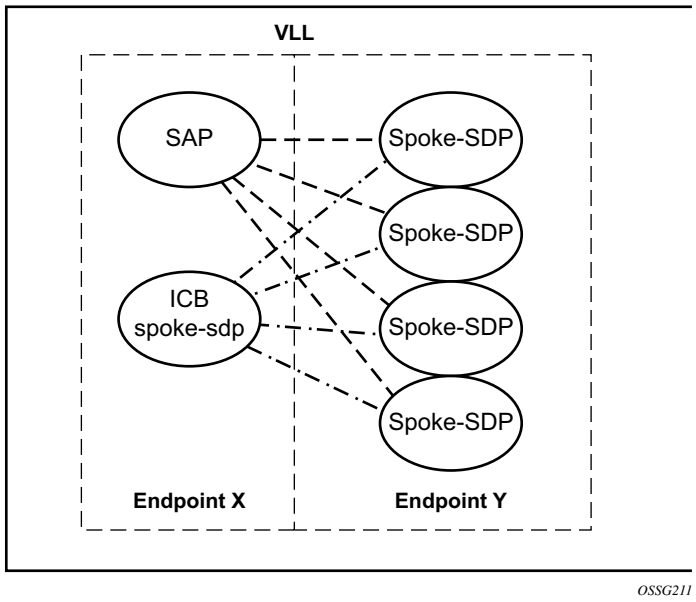

Figure 28: Redundant VLL Endpoint Objects

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. They are referred to as endpoint 'X' and endpoint 'Y' as illustrated in Figure 28.

Note that Figure 28 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.

T-LDP Status Notification Handling Rules

Referring to Figure 28 on page 173 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 173 can be used on endpoints “X” and “Y”. The following sections refer to the specific combination objects in Figure 28 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).

Processing and Merging

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

If the SAP in endpoint “X” transitions locally to the down state, or received a SAP down notification by SAP-specific OAM signal, the node must send T-LDP SAP down status bits on the “Y” endpoint ICB spoke SDP only. Note 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 by remote T-LDP status bits or by 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 by the SAP specific OAM signal if applicable. An Ethernet SAP does not support signaling status notifications.

Epipe Configuration for MPLS-TP

Note: MPLS-TP is supported only on 7210 SAS-T in network mode. It is not supported on 7210 SAS access-uplink mode and on 7210 SAS-M in network mode. MPLS-TP PWs are supported in Epipe service.

The following subsections describe how SDPs and spoke-SDPs are used with MPLS-TP LSPs and static PWs with MPLS-TP OAM.

SDPs

An SDP used for MPLS-TP supports the configuration of an MPLS-TP identifier as the far end address, as an alternative to an IP address. IP addresses are used if IP/MPLS LSPs are used by the SDP, or MPLS-TP tunnels identified by IPv4 source or destination addresses. MPLS-TP node identifiers are used if MPLS-TP tunnels are used.

The following CLI shows the new MPLS-TP options:.

```
config
   service
          sdp
             no description
             network-domain "default"
             signaling off
             far-end node-id 0.0.0.43 global-id 4294967295
             no mixed-lsp-mode
             no ldp
             no bgp-tunnel
             lsp "unnumberedLSP"
             no vlan-vc-etype
             no pbb-etype
             no path-mtu
             no adv-mtu-override
             keep-alive
                 shutdown
                 hello-time 10
                 hold-down-time 10
                 max-drop-count 3
                 timeout 5
                 no message-length
             exit
             no metric
             no collect-stats
             no accounting-policy
             binding
                 no port
             exit
             no shutdown
-------------------------------------------
*A:7210SAS>config>service>sdp#
```

The **far-end node-id** <ip-address> **global-id** <global-id> command is used to associate an SDP far end with an MPLS-TP tunnel whose far end address is an MPLS-TP node ID. If the SDP is associated with an RSVP-TE LSP, then the far-end must be a routable IPv4 address.

The system accepts the node-id being entered as either 4-octet IP address format <a.b.c.d> or unsigned integer format.

The SDP far-end refer to an MPLS-TP node-id or global-id only if:

- Delivery type is MPLS.
- Signalling is **off**.
- Keep-alive is disabled
- Mixed-lsp-mode is disabled
- Adv-mtu-override is disabled

An LSP will only be allowed to be configured if the far-end info matches the LSP far-end info (whether MPLS-TP or RSVP).

- Only one LSP is allowed if the far-end is an MPLS-TP node-id or global-id
- MPLS-TP or RSVP-TE LSPs are supported. However, note that LDP and BGP LSPs are not blocked in CLI.

Signaling TLDP or BGP is blocked if:

- Far-end node-id/global-id configured
- Control-channel-status enabled on any spoke (or mate vc-switched spoke)
- PW-path-id configured on any spoke (or mate vc-switched spoke)

The following commands are blocked if a far-end node-id or global-id is configured:

- Class-forwarding
- Tunnel-far-end
- Mixed-LSP-mode
- Keep-alive
- LDP or BGP-tunnel
- Adv-MTU-override

VLL Spoke SDP Configuration

7210 SAS-R6 can be S-PE or T-PE and 7210 SAS-T can only be a T-PE. MPLS-TP OAM related commands are applicable to spoke-sdps configured under all services supported by MPLS-TP

pseudowires. All commands and functions that are applicable to spoke-sdps in today's implementation are supported, except for those that explicitly depend on an LDP session on the SDP or as stated below. Likewise, all existing functions on a given service SAP are supported if the spoke-sdp that it is matched to is MPLS-TP.

The following describes how to configure MPLS-TP on an Epipe VLL. However, similar configuration applies to other VLL types.

A spoke-sdp bound to an SDP with the mpls-tp keyword cannot be **no shutdown** unless the ingress label, the egress label, the control word, and the pw-path-id are configured.

```
*7210SAS>config>service>epipe# info
----------------------------------------------
            sap 1/1/10:1.111 create
            exit
            spoke-sdp 1:111 create
                ingress
                    vc-label 2111
                exit
                egress
                    vc-label 2111
                exit
                control-word
                pw-path-id
                        agi 0:111
                        saii-type2 4294967295:0.0.0.42:111
                        taii-type2 4294967295:0.0.0.43:111
                exit
                no shutdown
            exit
            no shutdown
----------------------------------------------
*7210SAS>config>service>epipe#
```

The pw-path-id context is used to configure the end-to-end identifiers for a MS-PW. These may not coincide with those for the local node if the configuration is at an S-PE. The saii and taii are consistent with the source and destination of a label mapping message for a signaled PW.

The **control-channel-status** command enables static PW status signaling. This is valid for any spoke-sdp where **signaling none** is configured on the SDP (for example, where T-LDP signaling is not in use). The refresh timer is specified in seconds, from 10-65535, with a default of 0 (off). This value can only be changed if **control-channel-status** is **shutdown**. Commands that rely on PW status signaling are allowed if control-channel-status is configured for a spoke-sdp bound to an SDP with signaling off, but the system will use control channel status signaling rather than T-LDP status signaling. The ability to configure control channel status signaling on a given spoke-sdp is determined by the credit based algorithm described below. Control-channel-status for a particular PW only counts against the credit based algorithm if it is in a no shutdown state and has a non-zero refresh timer.

Note that a shutdown of a service will result in the static PW status bits for the corresponding PW being set.

The spoke-sdp is held down unless the **pw-path-id** is complete.

The system accepts the node-id of the pw-path-id saii or taii being entered as either 4-octet IP address format <a.b.c.d> or unsigned integer format.

The control-word must be enabled to use MPLS-TP on a spoke-sdp.

The **pw-path-id** only configurable if all of the following is true:

- Network mode D
- SDP signaling is off
- Control-word enabled (control-word is disabled by default)
- Service type epipe or VPLS
- Mate SDP signaling is off for vc-switched services
- An MPLS-TP node-id/global-id is configured under the **config>router>mpls>mpls-tp** context. This is required for OAM to provide a reply address.

In the vc-switching case, if configured on a mate spoke-sdp, then the TAII of the spoke-sdp must match the SAII of its mate, and SAII of spoke-sdp has to match the TAII of its mate.

A control-channel-status no shutdown is allowed only if all of the following is true:

- Network-mode D
- SDP signaling is off
- Control-word enabled (control-word by default is disabled)
- The service type is epipe or VPLS interface
- Mate SDP signaling is off (in vc-switched services)
- pw-status-signaling is enabled (see below)
- pw-path-id is configured for this spoke.

The **hash-label** option is only configurable if SDP far-end is not node-id or global-id.

The control channel status request mechanism is enabled when the request-timer <timer> parameter is non-zero. When enabled, this overrides the normal RFC-compliant refresh timer behavior. The refresh timer value in the status packet defined in RFC 6478 is always set to zero.

The refresh-timer in the sending node is taken from the request-timer. The two mechanisms are not compatible with each other. One node sends a request timer while the other is configured for refresh timer. In a given node, the request timer can only be configured with both acknowledgment and refresh timers disabled.

Once configured, the procedures below are used instead of the RFC 6478 procedures when a PW status changes.

The CLI commands to configure control channel status requests are shown, below:

```
[no] control-channel-status
[no] refresh-timer <value> //0,10-65535, default:0
[no] request-timer
[timeout-multiplier <value>]
[no] shutdown
exit
request-timer <timer1>: 0, 10-65535, defaults: 0.
```

- This parameter determines the interval at which PW status messages, including a reliable delivery TLV, with the “request” bit set (see below) are sent. This cannot be enabled if refresh-timer not equal to zero (0). retry-timer : 3-60s
- This parameter determines the timeout interval if no response to a PW status is received. This defaults to zero (0) when no retry-timer. timeout-multiplier <value> - 3-15.
- If a requesting node does not hear back after retry-timer times multiplier, then it must assume that the peer is down. This defaults to zero (0) when no retry-timer.

Credit Based Algorithm

In order to constrain the CPU resources consumed processing control channel status messages, the system should implement a credit-based mechanism. If a user enables control channel status on a PW[n], then a certain number of credits *c_n* are consumed from a CPM-wide pool of *max_credit* credits. The number of credits consumed is inversely proportional to the configured refresh timer (the first three messages at 1 second interval do not count against the credit). If the *current_credit <= 0*, then control channel status signaling cannot be configured on a PW (but the PW can still be configured and no shut).

The following is an example algorithm:

If *refresh timer > 0, c_n = 65535 / refresh_timer*

Else *c_n = 0.*

For *n=1, current_credit[n] = max-credits – c_n*

Else *current_credit [n] = current_credit [n-1] – c_n*

If a PE with a non-zero refresh timer configured does not receive control channel status refresh messages for 3.5 time the specified timer value, then by default it will time out and assume a PW status of zero. A proprietary optional extension to the [RFC6478] protocol should be implemented to enable a node to resolve such a stale PW status condition by requesting the status from the far end node in certain cases.

VLAN Range for SAPs in an Epipe Service

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

NOTE: Grouping a range of VLAN IDs to a SAP is supported only for Virtual Leased Lines (VLL) Ethernet services.

Processing behavior for SAPs using VLAN ranges in access-uplink mode

The access SAPs that specifies VLAN range values using connection-profile (also known as, dot1q range SAPs) is allowed in Epipe service and in VPLS service. For more information on functionality supported, see VLAN Range SAPs feature Support and Restrictions on page 182. The system allows only one range SAP in an Epipe service. It fails any attempt to configure more than one range SAP in an Epipe service. Range SAP can be configured only on access ports. The other endpoint in the Epipe service has to be a “Q.* SAP” in access-uplink mode. The processing and forwarding behavior for packets received on range SAPs are listed below:

- No VLAN tags are removed/stripped on ingress of access dot1q SAP configured to use VLAN ranges. A single tag (Q1) is added to the frame when it is forwarded out of the Q1.* access-uplink SAP.
- When a packet is received on the access-uplink Q1.* SAP, the outermost tag is removed and the packet is forwarded out of the access dot1q range SAP. The system does not check if the inner VLAN tag matches the VLANs IDs (both range and individual values specified in the “connection-profile”) of the dot1q access SAPs configured in the service.
- The dot1q range sap can be supported in a service with svc-sap-type set to ‘dot1q-range’.
- Support available for 7210 SAS-M and 7210 SAS-T in Access-Uplink mode.

VLAN Range SAPs feature Support and Restrictions

- The access SAPs that specifies VLAN range values (using connection-profile) is allowed only in E-Pipe service. The system allows only one range SAP in an Epipe service. It will fail any attempt to configure more than one range SAP in an Epipe service. Range SAP can be configured only on access ports.
- In access-uplink mode, the dot1q range sap is allowed to be configured only in a service with svc-sap-type set to ‘dot1q-range’. In network mode, the dot1q range sap is allowed to be configured in a service with svc-sap-type set to ‘any’.

- The access SAPs using VLAN range values are allowed only for Dot1q encapsulation port or LAG. A connection profile is used to specify either range of VLAN IDs or individual VLANs to be grouped together in a single SAP.
- A “connection profile” is used to specify either range of VLAN IDs or individual VLANs to be grouped together in a single SAP.
- No Dot1q default sap is allowed on the same access port as the one on which a SAP with a range is configured.
- Multiple “connection-profile” can be used per port or Lag as long as the VLAN value specified by each of them does not overlap. The number of VLAN ranges available per port/LAG is limited. The available number must be shared among all the SAPs on the port/LAG.
- “Connection-profile”, associated with a SAP cannot be modified. To modify a connection profile, it must be removed from all SAPs that are using it.

Processing behavior for SAPs using VLAN ranges in network mode

- The access SAPs that specifies VLAN range values (using connection-profile) is allowed only in an epipe service. The system allows only one range SAP in an Epipe service. It will fail any attempt to configure more than one range SAP in an Epipe service. Range SAP can be configured only on access ports. The other endpoint in the Epipe service has to be a Q.* access SAP or a spoke-sdp (PW) in network mode. The Spoke-SDP processing and forwarding behavior for packets received on range SAPs are listed below: No VLAN tags are removed/stripped on ingress of the access dot1q SAPs using VLAN range connection profile. When the other endpoint in the service is configured to be an Q1.* access SAP, 7210 adds another tag to the packet and forwards it out of that SAP. If the other endpoint in the service is configured to be a spoke-SDP whose vc-type is set to vc-ether, 7210 SAS adds the appropriate MPLS PW and LSP encapsulations and forwards it out of the SDP. In the reverse direction, when the other endpoint is a Q1.* SAP and a packet is received on it, 7210 SAS removes the outermost VLAN tag and forwards the packet out of the access dot1q SAP using VLAN ranges. When the other endpoint is a spoke-sdp (whose vc-type is set to vc-ether), 7210 SAS removes the MPLS PW and LSP encapsulation and forwards the packet out of the access dot1q SAP using VLAN ranges. The system does not check if the VLAN in the packet matches the VLAN IDs of the dot1q access SAPs configured in the service.
- ACL support - Filter policies are supported on SAP ingress. In 7210 SAS-M and 7210 SAS-T access-uplink mode, IP criteria and MAC criteria based filter policy is supported with access SAPs. **Note**: For more information on ACL on range SAPs, see The 7210 SAS Router Configuration Guide.
- In 7210 SAS-M and 7210 SAS-T network mode, only MAC criteria based filter policy supported with access SAPs. **Note**: For more information on ACL on range SAPs, see The 7210 SAS Router Configuration Guide.
- QoS – Ingress classification, metering with hierarchical metering, marking, queuing and shaping for SAP ingress and SAP egress. On egress per port queues and shaping is available on 7210 SAS-M.
  - → SAP ingress classification criteria is available for use with VLAN range SAPs is similar to that available for other SAPs supported in an Epipe service. Dot1p based ingress classification uses the Dot1p bits in the outermost VLAN tag for matching. On access egress, dot1p received from the SDP (on a network port) from another access port is preserved.
- The amount of hardware resources (such as CAM entries used for matching in QoS classification and ACL match, meters used in SAP ingress policy, and others.) consumed by a single range SAP is equivalent to the amount of resources consumed by a single SAP that specifies a single VLAN ID for service identification. In other words, the hardware has the ability to match a range of VLAN values and hence uses ‘X’ resources for a SAP using a VLAN range instead of X * n, where ‘n’ is the number of VLANs specified in the range and X is the amount of QoS or ACL resources needed.

- Ingress accounting support is similar to the support available for other SAPs in an Epipe service. Count of packets or octets received from individual VLANs configured in the connection profile is not available. No support for Egress SAP statistics and accounting is available.
- Mirroring is supported. In network mode, the use of service resiliency mechanisms such as MC-LAG and Epipe PW redundancy is supported.

VLL Service Considerations

This section describes various of the general 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 routers.
- An SDP encapsulation type, MPLS.

SAP Encapsulations

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

- Ethernet null
- Ethernet dot1q
- QinQ

Note that while different encapsulation types can be used, encapsulation mismatch 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 potentially be double tagged when it is transmitted out of the Dot1q SAP.

QoS Policies

Traffic Management - Traffic management of Ethernet VLLs is achieved through the application of ingress QoS policies to SAPs and access egress QoS policies applied to the port. All traffic management is forwarding-class aware and the SAP ingress QoS policy identifies the forwarding

class based on the rules configured to isolate and match the traffic ingressing on the SAP. Forwarding classes are determined based on the Layer 2 (Dot1p, MAC) or Layer 3 (IP, DSCP) fields of contained packets and this association of forwarding class at the ingress will determine both the queuing and the Dot1P bit setting of packets on the Ethernet VLL on the egress.

SAP ingress classification and Policing - The traffic at the SAP ingress is classified and metered according to the SLA parameters. All the traffic ingressing on the SAP is classified to a particular forwarding class. All the forwarding class is metered through and marked in-profile or put-profile based on the Meter parameters.

When applied to 7210 SAS M Epipe services, service ingress QoS policies only create the unicast queues defined in the policy. The multipoint queues are not created on the service. Note that both Layer 2 or Layer 3 criteria can be used in the QoS policies for traffic classification in a service.

Egress Network DOT1P Marking - Marking of IEEE DOT1P bits in VLAN tag is as per the FC-to-Dot1p map. For details see the default network QoS policy in the QoS user guide. This marking is applied at the port level on access ports and access uplink ports.

Ingress Network Classification - Ingress network classification is based on the Dot1p bits in the outer VLAN tag received on the access uplink port. Dot1p-to-FC mapping is based on the network ingress QoS policy.

Filter Policies

7210 SAS Epipe services can have a single filter policy associated on both ingress and egress. Both MAC and IP filter policies can be used on Epipe services.

MAC Resources

Epipe services are point-to-point layer 2 VPNs capable of carrying any Ethernet payloads. Although an Epipe is a Layer 2 service, the 7210 SAS M 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:

Basic Configurations

Common Configuration Tasks

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

- Associate the service with a customer ID.
- Define SAP parameters
  - → Optional - select ingress QoS policies (configured in the **config>qos** context).
  - → Optional - select accounting policy (configured in the **config>log** context).
- Define spoke SDP parameters (Not applicable for 7210 SAS devices configured in Access Uplink mode).
- Enable the service.

Configuring VLL Components

This section provides VLL configuration examples for the VLL services:

Creating a Cpipe Service

Basic Configuration

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

- Customer ID
- Interface parameters
- Spoke SDP parameters

The following example displays a sample configuration of a Cpipe service.

```
*A:ces-A>config>service# cpipe 1
*A:ces-A>config>service>cpipe# info
----------------------------------------------
            sap 1/2/1.1 create
            exit
            spoke-sdp 12:1 create
            exit
            no shutdown
----------------------------------------------
```

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

CLI Syntax: `config>service# cpipe service-id [customer customer-id] [vpn vpn-id] [vc-type {satop-e1 | satop-t1 | cesopsn | cesopsn-cas}]`

The following displays a Cpipe service configuration example.

```
*A:ces-A>config>service>cpipe#
        cpipe 1 customer 1 vc-type satop-t1 create
            sap 1/2/1.1 create
                ingress
                    qos 12
                exit
            exit
            spoke-sdp 12:1 create
            exit
            no shutdown
        exit
*A:ces-A>config>service>cpipe# exit all

*A:ces-A>config>service>cpipe# info detail
----------------------------------------------
            no description
            service-mtu 1514
            sap 1/2/1.1 create
                no description
                cem
                    packet jitter-buffer 5 payload-size 192
                    report-alarm stray malformed pktloss overrun underrun
```

```
                    no report-alarm rpktloss rfault rrdi
                    no rtp-header
                exit
                ingress
                    qos 1
                exit
                no collect-stats
                no accounting-policy
                no shutdown
            exit
            spoke-sdp 12:1 create
                ingress
                    no vc-label
                exit
                egress
                    no vc-label
                exit
                no collect-stats
                no accounting-policy
                no precedence
                no shutdown
            exit
            no shutdown
-------------------------------------------
*A:Dut-A>config>service>cpipe# info detail
-------------------------------------------
            no description
            service-mtu 1514
            endpoint "y" create
                no active-hold-delay
                no description
                no revert-time
            exit
            sap 1/2/1.2 create
                no description
                cem
                    packet jitter-buffer 32 payload-size 64
                    report-alarm stray malformed pktloss overrun underrun
                    no report-alarm rpktloss rfault rrdi
                    no rtp-header
                exit
                ingress
                    qos 1
                    no aggregate-meter-rate
                exit
                no collect-stats
                no accounting-policy
                no shutdown
            exit
            spoke-sdp 123:104 endpoint "y" create
                ingress
                    no vc-label
                exit
                egress
                    no vc-label
                exit
                no collect-stats
                no accounting-policy
                no precedence
                no shutdown
            exit
            no shutdown
-------------------------------------------
```

```
*A:Dut-A>config>service>cpipe#
```

Configuration Requirements

Before a Cpipe service can be provisioned, the following tasks must be completed:

- Configuring a DS1 Port on page 193
- Configuring a Channel Group on page 194

Configuring a DS1 Port

The following displays an example of a DS1 port configured for CES.

```
*A:ces-A# configure port 1/2/1
*A:ces-A>config>port# info
----------------------------------------------
        tdm
            ds1
                framing ds1-unframed
                clock-source adaptive
                report-alarm looped
                channel-group 1
                    no shutdown
                exit
                no shutdown
            exit
        exit
        no shutdown
----------------------------------------------
*A:ces-A>config>port#
```

Configuring a Channel Group

The following displays an example of a DS1 channel group configured for CES.

```
*A:ces-A>config>port# info
----------------------------------------------
        tdm
            ds1
                framing ds1-unframed
                clock-source adaptive
                report-alarm looped
                channel-group 1
                    no shutdown
                exit
                no shutdown
            exit
        exit
        no shutdown
----------------------------------------------
*A:ces-A>config>port#


*A:ces-A>config>port# info detail
----------------------------------------------
        description "DS1/E1"
        tdm
            buildout short
            length 133
            ds1
                framing ds1-unframed
                no loopback
                clock-source adaptive
                report-alarm ais los
                no report-alarm oof rai looped
                channel-group 1
                    description "DS0GRP"
                    mode access
                    encap-type cem
                    timeslots 1-24
                    idle-payload-fill all-ones
                    no shutdown
                exit
                no shutdown
            exit
            line-impedance 100
        exit
        no shutdown
----------------------------------------------
```

Configuring Cpipe SAPs and Spoke SDPs

The following output displays examples of Cpipe SAP and spoke SDP configurations.

```
*A:ces-A>config>service>cpipe# info
----------------------------------------------
        cpipe 1 customer 1 vc-type satop-t1 create
            sap 1/2/1.1 create
                ingress
                    qos 12
                exit
            exit
            spoke-sdp 12:1 create
            exit
            no shutdown
        exit
----------------------------------------------
*A:ces-A>config>service>cpipe#
```

Creating an Epipe Service in Network Mode

Use the following CLI syntax to create an Epipe service.

CLI Syntax: `config>service# epipe service-id [customer customer-id] [create] [vpn vpn-id]`
`description description-string`
`no shutdown`

The following displays an Epipe configuration example:

```
A:ALA-1>config>service# info
----------------------------------------------
...
        epipe 1101 customer 1 vpn 1101 create
            description "Default epipe description for service id 1101"
            no shutdown
        exit
----------------------------------------------
A:ALA-1>config>service#
```

Creating an Epipe Service (for 7210 SAS-M and 7210 SAS-T in access uplink mode)

Use the following CLI syntax to create an Epipe service:

CLI Syntax: `config>service# epipe service-id [customer customer-id] [create] [svc-sap-type {null-star | dot1q | dot1q-preserve | any}] [customer- vid vlan-id] description description-string no shutdown`

```
A:ALA-1>config>service# info
----------------------------------------------
...
     epipe 500 customer 1 svc-sap-type null-star create
          description "Local Epipe Service with NULL SVC_SAP_TYPE"
          no shutdown
     exit
----------------------------------------------
A:ALA-1>config>service#
```

Configuring Epipe SAP Parameters

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

Use the following CLI syntax to create:

- Local Epipe SAPs on page 198
- Distributed Epipe Service on page 200

CLI Syntax:

```
config>service# epipe service-id [customer customer-id]
        sap sap-id
              accounting-policy policy-id
              collect-stats
              description description-string
              no shutdown
              egress
                 filter {ip ip-filter-name | mac mac-filter-name}
              ingress
                 filter {ip ip-filter-name | mac mac-filter-name}
                 qos policy-id
```

Local Epipe SAPs

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

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

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

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

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

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

Creating an Epipe Service for 7210 SAS-M with range SAPs

The following displays an example of connection-profile used to configure a range of SAPs and an Epipe configuration using the connection profile:

```
*A:7210SAS>config>connprof# info
```

```
-----------------------------------------------
        ethernet
            ranges 0 2804-2805 2810-2811 2813 2832-2839
        exit
-----------------------------------------------

*A:7210SAS>config>service>epipe# info
-----------------------------------------------
            description "Default epipe description for service id 292"
            sap 1/1/4:292.* create
                description "Default sap description for service id 292"
                exit
            exit
            sap 1/1/9:cp-292 create
                description "Default sap description for service id 292"
                exit
            exit
            no shutdown
-----------------------------------------------
```

Distributed Epipe Service

Note: SDPs are not supported by 7210 SAS devices configured in Access Uplink mode.

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

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

Meters (defined in sap-ingress policies) can be applied on ingress. It is associated with SAPs. Scheduler Policies can be applied on egress. It is associated with a port.

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

For SDP configuration information, see Configuring an SDP on page 71. For SDP binding information, see Configuring SDP Bindings on page 205.

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

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

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

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

```
A:ALA-1>config>service# info
----------------------------------------------
...
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to east coast"
            sap 221/1/3:21 create
                ingress
                    qos 555
                    filter ip 1
```

```
                exit

            exit
        exit
...
-----------------------------------------------
A:ALA-1>config>service#


A:ALA-2>config>service# info
-----------------------------------------------
...
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to west coast"
            sap 441/1/4:550 create
                ingress
                    qos 654
                    filter ip 1020
                exit

            exit
        exit
...
-----------------------------------------------
A:ALA-2>config>service#
```

Configuring Ingress SAP Parameters

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

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

This example displays SAP ingress and egress parameters.

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

The following example displays the Epipe SAP ingress configuration:

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

Configuring Default QinQ SAPs for Epipe Transit Traffic in a Ring Scenario

Note: Default QinQ SAPs are supported only on 7210 SAS devices configured in access-uplink mode.

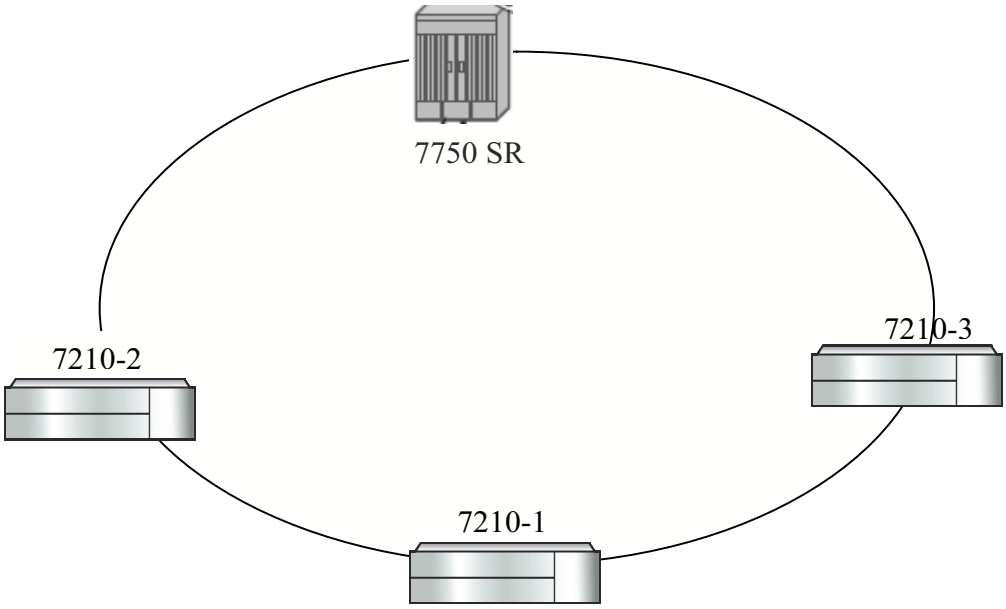

Figure 29: Default QinQ SAP for Transit Traffic in a Ring Scenario

In the Figure 29, 7210-1 is used to deliver some services to customers connected to the device and additionally it needs to pass through transit from other nodes on the ring (example – traffic from 7210-2 to 7210-3 OR from 7210-2 to 7750 –SR onto the core network).

Without Default QinQ SAPs, user would need to configure a service on 7210-1, with access-uplink SAPs for each service originating on some other node in the ring. With support for Default QinQ SAPs, all traffic which does not need to be delivered to any customer service configured on 7210-1 can be switched using the EPIPE service. The example shown below provides the sample configuration commands in this scenario:

```
ALA-1>config>service# epipe 8 customer 1 svc-sap-type null-star create
            sap 1/1/5:*.* create
                statistics
                    ingress
                    exit
                exit
            exit
            sap 1/1/6:*.* create
                statistics
                    ingress
                    exit
                exit
            exit
            no shutdown
```

```
exit
```

Configuring SDP Bindings

Note: SDPs are not supported by 7210 SAS-M and 7210 SAS-T devices configured in Access Uplink mode.

Figure 30 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.

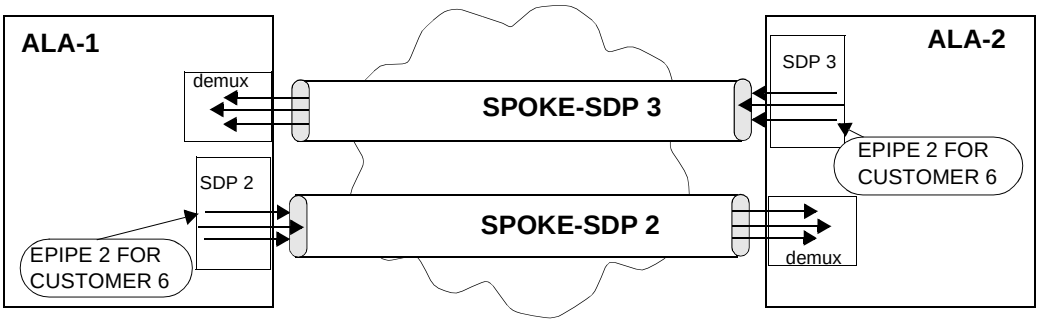

Figure 30: SDPs — Uni-Directional Tunnels

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 Service on page 200).

A:ALA-1>config>service# epipe 5500

```
    config>service>epipe# spoke-sdp 2:123
    config>service>epipe>spoke-sdp# egress
    config>service>epipe>spoke-sdp>egress# vc-label 5500
    config>service>epipe>spoke-sdp>egress# exit
    config>service>epipe>spoke-sdp# ingress
    config>service>epipe>spoke-sdp>ingress# vc-label 6600
    config>service>epipe>spoke-sdp>ingress# exit
    config>service>epipe>spoke-sdp# no shutdown

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

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

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


A:ALA-2>config>service# info
----------------------------------------------
...
exit
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to west coast"
```

```
            sap 441/1/4:550 create
                ingress
                    qos 654
                    filter ip 1020
                exit
            exit
            spoke-sdp 2:456 create
                ingress
                    vc-label 5500
                exit
                egress
                    vc-label 6600
                exit
            exit
            no shutdown
        exit
...
-------------------------------------------
A:ALA-2>config>service#
```

Using Spoke SDP Control Words

Note: SDPs are not supported by 7210 SAS devices configured in Access Uplink mode.

The control word command provides the option to add a control word as part of the packet encapsulation for PW types for which the control word is optional. On 7210 an option is provided to enable it for Ethernet PW (Epipe). The control word might be needed because when ECMP is enabled on the network, packets of a given PW may be spread over multiple ECMP paths if the hashing router mistakes the PW packet payload for an IPv4 or IPv6 packet. This occurs when the first nibble following the service label corresponds to a value of 4 or 6.

The control word negotiation procedures described in Section 6.2 of RFC 4447 are not supported and therefore the service will only come up if the same C bit value is signaled in both directions. If a spoke-sdp is configured to use the control word but the node receives a label mapping message with a C-bit clear, the node releases the label with an “Illegal C-bit” status code per Section 6.1 of RFC 4447. As soon as the user enables control of the remote peer, the remote peer withdraws its original label and sends a label mapping with the C-bit set to 1 and the VLL service is up in both nodes.

When the control word is enabled, VCCV packets also include the VCCV control word. In that case, the VCCV CC type 1 (OAM CW) is signaled in the VCCV parameter in the FEC. If the control word is disabled on the spoke-sdp, then the Router Alert label is used. In that case, VCCV CC type 2 is signaled. Note that for a multi-segment PW (MS-PW), the CC type 1 is the only supported and thus the control word must be enabled on the spoke-sdp to be able to use VCCV-ping and VCCV-trace.

The following displays a spoke SDP control word configuration example:

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

Configuring VLL Resilience

Figure 31 displays 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.

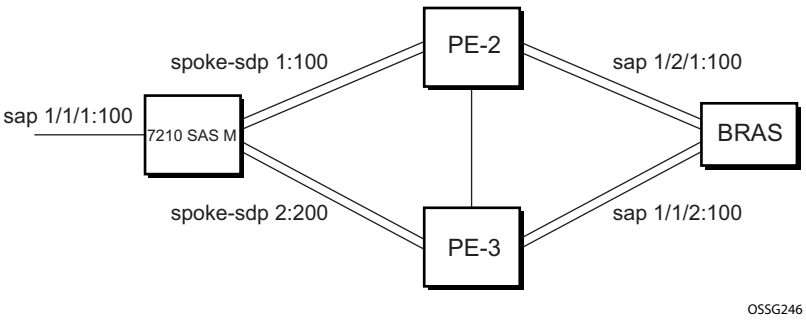

Figure 31: VLL Resilience

PE1:

The following displays an example for the configuration on PE1.

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

Configuring VLL Resilience for a Switched Pseudowire Path

Note that the 7210 SAS M only supports T-PE functionality.

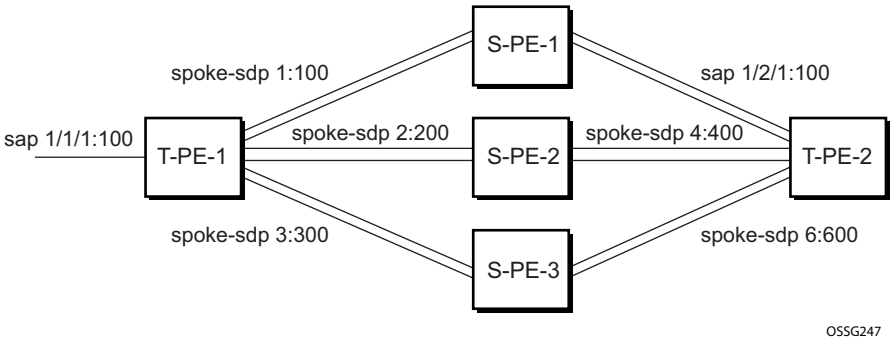

Figure 32: VLL Resilience with Pseudowire Switching

T-PE1

The following displays an example for the configuration on TPE1.

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

T-PE2

The following displays an example for the configuration on TPE2.

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

S-PE1

The following displays an example for the configuration on S-PE1.

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

Service Management Tasks

This section discusses the following Cpipe service management tasks:

This section discusses the following Epipe service management tasks:

Modifying a Cpipe Service

The following example displays the Cpipe service configuration.

```
*A:ces-A>config>service>cpipe# info
----------------------------------------------
        cpipe 1 customer 1 vc-type satop-t1 create
            sap 1/2/1.1 create
                ingress
                    qos 12
                exit
            exit
            spoke-sdp 12:1 create
            exit
            no shutdown
        exit
----------------------------------------------
*A:ces-A>config>service>cpipe
```

Deleting a Cpipe Service

A Cpipe service cannot be deleted until SAPs are shut down and deleted. If a spoke-SDP is defined, it must be shut down and removed from the configuration as well.

Use the following CLI syntax to delete a Cpipe service:

CLI Syntax:

```
config>service#
        [no] cpipe service-id [customer customer-id]
           [no] spoke-sdp sdp-id
              [no] shutdown
           shutdown
```

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 1/1/3:21
        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 1/1/3:21 create
               accounting-policy 14
             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.
2. Delete the SAP.
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
```

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

VLL Services Command Reference

Command Hierarchies

Cpipe Service Configuration Commands

```
config
    — service
        — cpipe service-id [customer customer-id] [vpn vpn-id] [vc-type {satop-e1 | satop-t1 | cesopsn
          | cesopsn-cas}] [create]
        — no cpipe service-id
            — description description-string
            — no description [description-string]
            — endpoint endpoint-name [create]
            — no endpoint endpoint-name
                — active-hold-delay active-endpoint-delay
                — no active-hold-delay
                — description description-string
                — no description [description-string]
                — revert-time revert-time infinite
                — no revert-time
            — sap sap-id [no-endpoint] [create]
            — sap sap-id endpoint endpoint-name [create]
            — no sap sap-id
            — [no] service-name
                — accounting-policy acct-policy-id
                — no accounting-policy [acct-policy-id]
                — cem
                    — packet jitter-buffer milliseconds [payload-size bytes]
                    — packet payload-size bytes
                    — no packet
                    — [no] report-alarm [stray] [malformed] [pktloss] [overrun]
                      [underrun] [rpktloss] [rfault] [rrdi]
                    — [no] rtp-header
                — [no] collect-stats
                — description description-string
                — no description [description-string]
                — ingress
                    — [no] qos [policy-id]
            — service-mtu octets
            — no service-mtu
            — [no] service-name
            — [no] shutdown
            — spoke-sdp sdp-id[:vc-id] [no-endpoint] [create]
            — spoke-sdp sdp-id:vc-id [create] endpoint endpoint-name [icb]
            — no spoke-sdp sdp-id[:vc-id]
                — accounting-policy acct-policy-id
                — no accounting-policy
```

— **description** *description-string*
— **no description**
— [**no**] **collect-stats**
— **egress**
— **vc-label** *egress-vc-label*
— **no vc-label** [***egress-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

Epipe Global Commands

config
— **service**
— [**no**] **epipe** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [svc-sap-type { any | qinq-inner-tag-preserve}] (for 7210 SAS devices in Network mode)
— **epipe** *service-id* [**customer** *customer-id*] [**create**][**vpn** *vpn-id*] [**svc-sap-type** {**null-star** | **dot1q** | **dot1q-preserve**|**any**| **qinqinner-tag-preserve**}](for 7210 SAS devices in Access uplink mode)
— **no epipe** *service-id*
— **description** *description-string*
— **no description**
— [**no**] **endpoint** *endpoint-name* **[create]**
— **active-hold-delay** *active-endpoint-delay*
— **no active-hold-delay**
— **revert-time** [*revert-time* | **infinite**]
— **no revert-time**
— **standby-signaling-master**
— [**no**] **standby-signaling-master**
— **sap** *sap-id* [**create**]
— **no sap** *sap-id*
— **service-mtu** *octets* (for 7210 SAS-M and 7210 SAS-T in Network mode)
— **no service-mtu**
— [**no**] **service-mtu-check** (for 7210 SAS-M and 7210 SAS-T in Network mode)

- [**no**] **shutdown**
- **spoke-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] [**no-endpoint**]
- **spoke-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] **endpoint**
- **no spoke-sdp** *sdp-id*[:*vc-id*]

Epipe SAP Configuration Commands

```
config
    — service
        — epipe service-id [customer customer-id] [create] [vpn vpn-id][svc-sap-type { any | qinq-inner-tag-preserve} ]  (for 7210 SAS-M and 7210 SAS-T in Network mode)
        — epipe service-id [customer customer-id] [create] [vpn vpn-id][customer customer-id] [create] [vpn vpn-id] [svc-sap-type {null-star| dot1q| dot1q-preserve| any| dot1q-range}] [customer-vid vlan-id] (for 7210 SAS-M and 7210 SAS-T in Access uplink mode)
            — sap sap-id [no-endpoint] [create] with-aggregate-meter
            — sap sap-id [endpoint endpoint-name]  [create]
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy acct-policy-id
                — [no] collect-stats
                — description description-string
                — no description
                — eth-cfm
                    — [no] mep mep-id domain md-index association ma-index [direction {up | down}]
                        — [no] ais-enable
                            — [no] client-meg-level [[level [level ...]]
                            — [no] interval {1 | 60}
                            — [no] priority priority-value
                        — [no] ccm-enable
                        — [no] ccm-ltm-priority priority
                        — [no] description
                        — [no] eth-test-enable
                            — [no] bit-error-threshold bit-errors
                            — [no] test-pattern {all-zeros | all-ones} [crc-enable]
                        — [no] fault-propagation-enable {use-if-tlv | suspend-ccm}
                        — low-priority-defect {allDef | macRemErrXcon | remErrXcon | errXcon | xcon | noXcon}
                        — [no] mac-address mac-address
                        — [no] one-way-delay-threshold seconds
                        — [no] shutdown
                    — mip [mac mac address]
                    — mip default-mac
                    — no mip
                    — egress
                        — filter [ip ip-filter-id]
                        — filter  [ ipv6 ipv6 -filter-id]
                        — filter [mac mac-filter-id] (app
                        — no filter [ip ip-filter-id] [ ipv6 ipv6 -filter-id] [mac mac-filter-id]
                    — ingress
                        — aggregate-meter-rate <rate-in-kbps> [burst <burst-in-kbits>]
                        — no aggregate-meter-rate
                        — filter [ip ip-filter-id]
                        — filter [ ipv6 ipv6-filter-id]
                        — filter [mac mac-filter-id]
                        — no filter [ip ip-filter-id] [ ipv6 ipv6-filter-id] [mac mac-filter-id]
```

```
                    — meter-override
                    — no meter-override
                        — meter meter-id [create]
                        — no meter meter-id
                            — adaptation-rule [pir adaptation-rule] [cir
                              adaptation-rule]
                            — cbs size-in-kbytes
                            — no cbs
                            — mbs size-in-kbits
                            — no mbs
                            — mbs mode
                            — no mode
                            — no mode
                            — rate cir cir-rate [pir pir-rate]
                    — qos policy-id
                    — no qos
— ethernet
        — [no] llf
— [no] ignore-oper-down
— [no] shutdown
— statistics
        — ingress
                    — counter-mode {in-out-profile-count|forward-drop-
                      count}
                    — [no] shutdown
— tod-suite tod-suite-name
— no tod-suite
```

Epipe Spoke SDP Configuration Commands

Note: Spoke SDP commands are not supported on 7210 SAS devices configured in Access Uplink mode. It is supported only on network mode.

config
- **service**
  - **epipe** *service-id* **[customer** *customer-id***] [create] [vpn** *vpn-id***] [svc-sap-type {any|qinq-inner-tag-preserve}] [customer-vid** *vlan-id***] [pbb-epipe]** (for 7210 SAS-M and 7210 SAS-T in Network mode)
  - **epipe** *service-id* **[customer** *customer-id***] [create] [svc-sap-type {null-star|dot1q-pre-serve|any|dot1q-range}] [customer-vid** *vlan-id***]** (for 7210 SAS-M and 7210 SAS-T in Access uplink mode)
    - **spoke-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] [**no-endpoint**]
    - **spoke-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] **endpoint**
    - **no spoke-sdp** *sdp-id*[:*vc-id*]
      - **accounting-policy** *acct-policy-id*
      - **no accounting-policy**
      - [**no**] **collect-stats**
      - [**no**] **control-word**
      - **control-channel-status**
        - **acknowledgment**
        - **no acknowledgment**
        - **refresh-timer** ***seconds***
        - **no refresh-timer**
        - **request-timer request-timer** *request-timer-secs* **retry-timer** *retry-timer-secs* **timeout-multiplier** *multiplier*
      - [**no**] **description**
      - [**no**] **egress**
        - [**no**] **vc-label** *egress-vc-label*
      - **eth-cfm**
        - [**no**] **ais-enable**
        - [**no**] **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {up | down}]
          - [**no**] **ais-enable**
            - [**no**] **client-meg-level** [[**level** [level ...]]
            - [**no**] **interval** {**1** | **60**}{**1** | **60**}
            - [**no**] **priority** *priority-value*
          - [**no**] **ccm-enable**
          - [**no**] **ccm-ltm-priority** *priority*
          - [**no**] **description**
          - [**no**] **eth-test-enable**
            - [**no**] **bit-error-threshold** *bit-errors*
            - [**no**] **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
          - **[no] fault-propagation-enable {use-if-tlv | suspend-ccm}**
          - **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}
          - [**no**] **mac-address** *mac-address*
          - [**no**] **one-way-delay-threshold** *seconds*
          - [**no**] **shutdown**
        - **mip [mac** *mac address***]**
        - **mip default-mac**
        - **no mip**

- [**no**] **force-vlan-vc-forwarding**
- [**no**] **ingress**
  - [**no**] **vc-label** *egress-vc-label*
- **precedence** [*precedence-value*| **primary**]
- **no precedence**
- **[no] pw-path-id**
  - **agi** *attachment-group-identifier*
  - **no agi**
  - **no saii-type2**
  - **saii-type2** *global-id:node-id:ac-id*
  - **no taii-type2**
  - **taii-type2** *global-id:node-id:ac-id*
- **no pw-status-signaling**
- **pw-status-signaling**
- [**no**] **shutdown**
- **vlan-vc-tag** *0..4094*
- **no vlan-vc-tag** [*0..4094*]

Connection Profile Commands

config

- **connection-profile** *conn-prof-id [create]*
- **no connection-profile** *conn-prof-id*
  - **description** *description-string*
  - **no description**
  - **ethernet**
    - **no ranges**
  - **ranges** *vlan ranges [vlan ranges...(upto 32 max)]*

Show Commands

```
show
    — service
        — egress-label start-label [end-label]
        — id service-id
            — all
            — base
            — endpoint [endpoint-name]
            — epipe
            — labels
            — stp [sap-id] [detail]]
        — sap-using [sap sap-id]
        — sap-using [ingress | egress] filter filter-id
        — sap-using [ingress ] qos-policy qos-policy-id
        — service-using [epipe] [cpipe] [vpls] [mirror] [cpipe] [b-vpls] [i-vpls] [m-vpls] [sdp sdp-id]
          [customer customer-id]
show
    — connection-profile [conn-prof-id] [associations]
```

Clear Commands

```
clear
    — service
        — id service-id
        — statistics
            — id service-id
            — counters
            — sap sap-id {all |cem | counters | stp| l2pt}
```

VLL Service Configuration Commands

Generic Commands

shutdown

Syntax [**no**] **shutdown**
config>service>cpipe
config>service>cpipe>sap
config>service>cpipe>spoke-sdp
config>service>epipe
config>service>epipe>sap
config>service>epipe>spoke-sdp
config>service>epipe>sap>eth-cfm>mep

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.

description

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

Context config>service>cpipe
config>service>cpipe>endpoint
config>service>cpipe>sap
config>service>epipe
config>service>epipe>sap
config>service>epipe>spoke-sdp
config>connection-profile

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.

eth-cfm

Syntax **eth-cfm**

Context config>service>vpls
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>vll>sap

Description This command enables the context to configure ETH-CFM parameters.

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {up | down}]
no mep *mep-id* **domain** *md-index* **association** *ma-index*

Context config>service>epipe>sap>eth-cfm

Description This command provisions the maintenance endpoint (MEP).

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

Note: For more information on ETH-CFM support for different services, see Table 10, “ETH-CFM Support Matrix for 7210 SAS-M Devices,” on page 99 Table 12, “ETH-CFM Support Matrix for 7210 SAS-D Devices,” on page 100.

Parameters *mep-id* — Specifies the maintenance association end point identifier.

Values 1 — 8191

md-index — Specifies the maintenance domain (MD) index value.

Values 1 — 4294967295

ma-index — Specifies the MA index value.

Values 1 — 4294967295

direction up| down — Indicates the direction in which the maintenance association (MEP) faces on the bridge port. Direction is not supported when a MEP is created directly under the vpls>eth-cfm construct (vMEP).

down — Sends ETH-CFM messages away from the MAC relay entity.

up — Sends ETH-CFM messages towards the MAC relay entity.

VLL Global Commands

cpipe

Syntax **cpipe** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*] [**vc-type** {**satop-e1** | **satop-t1** | **cesopsn** | **cesopsn-cas**}] [**create**]
no cpipe *service-id*

Context config>service

Description This command configures a Circuit Emulation Services instance. 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 services exist until they are explicitly created with this command.

The **no** form of this command deletes the service instance with the specified *service-id*. The service cannot be deleted until the service has been shutdown.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7210 SAS on which this service is defined.

Values *service-id*: 1 — 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-type — The vc-type defines the type of unstructured or structured circuit emulation service to be configured.

Values **satop-e1**: unstructured E1 circuit emulation service
satop-t1: unstructured DS1 circuit emulation service
cesopsn: basic structured n*64 kbps circuit emulation service
c**esopsn-cas:** structured n*64 kbps circuit emulation service with signaling

Default cesopsn

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

epipe

Syntax [no] **epipe** *service-id* [**customer** *customer-id*] [create] [**vpn** *vpn-id*] [**svc-sap-type** { any | qinq-inner-tag-preserve}] (for 7210 SAS devices in Network mode)
epipe *service-id* [**customer** *customer-id*] [create] [***vpn*** vpn-id] [**svc-sap-type** {null-star | dot1q | dot1q-preserve|any| qinqinner-tag-preserve}] [customer-vid vlan-id] (for 7210 SAS devices in Access uplink mode)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 7210 SAS.

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.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7210 SAS on which this service is defined.

Values *service-id*: 1 — 2147483648
svc-name: 64 characters maximum

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

svc-sap-type — Specifies the type of service and allowed SAPs in the service.

null-star — Specifies that the allowed SAP in the service that can be Null SAP, dot1q Default SAP, Q.* SAP or Default QinQ SAP (also known as *.* SAP). Supported only on 7210 SAS-M and 7210 SAS-T access-uplink mode.

dot1q — Specifies that the allowed SAP in the service are Dot1q SAPs and dot1q explicit null SAPs. Supported only in 7210 SAS-M and 7210 SAS-T access-uplink mode.

dot1q-preserve — Specifies that the allowed SAP in the service are Dot1q. The Dot1q ID is not stripped after packets matches the SAP. Supported only in 7210 SAS-M and 7210 SAS-T access-uplink mode.

dot1q-range — Specifies that the access SAP in the service can use VLAN ranges as the SAP tags. The VLAN ranges are configured using the configure> connection-profile CLI command. On ingress of the access dot1q SAP using VLAN ranges, the outermost tag is not removed before forwarding. Supported in 7210 SAS-M and 7210 SAS-T access-uplink mode.

any — When svc-sap-type is set to **any**, for a NULL SAP, the system processes and forwards only packets with no VLAN tag (that is, untagged). All other packets with one or more VLAN tags (even those with priority tag only) are not processed and dropped. Users can use the service with svc- sap-type set to **null-star**, to process and forward packets with one or more tags (including priority tag) on a null SAP.

qinq-inner-tag-preserve — When svc-sap-type is set to this value, an Epipe service processes and forwards packets received with 3 tags on a QinQ SAP. Please read the Epipe chapter above to learn more about the support available and restrictions that apply. Supported only in 7210 SAS-M and 7210 SAS-T in network mode.

Default null-star

customer-vid *vlan-id* **—** Defines the dot1q VLAN ID to be specified while creating the local Dot1q SAP for **svc-sap-type dot1q-preserve**.

Values 1 — 4094

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

endpoint

Syntax [**no**] **endpoint** *endpoint-name*

Context config>service>cpipe
config>service>epipe

Description This command configures a service endpoint.

Parameters *endpoint-name* — Specifies an endpoint name.

active-hold-delay

Syntax **active-hold-delay** *active-hold-delay*

no active-hold-delay

Context config>service>epipe>endpoint

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>epipe>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.

standby-signaling-master

Syntax [**no**] **standby-signaling-master**

Context config>service>vll>endpoint

Description When this command is enabled, the pseudowire standby bit (value 0x00000020) will be sent to T-LDP peer for each spoke-sdp of the endpoint that is selected as a standby.

This command is mutually exclusive with a VLL mate SAP created on a mc-lag/mc-aps or ICB. It is also mutually exclusive with vc-switching.

Default no standby-signaling-master

service-mtu

Syntax **service-mtu** *octets*
no service-mtu

Context config>service>epipe

Description This command configures the service payload (Maximum Transmission Unit – MTU), in bytes, for a service. The specified 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 operational states of SAP and SDP bindings in a service.

On 7210 SAS nodes, when a packet is received on a SAP, the service MTU check includes the length of the packet and the SAP delineation encapsulation overhead (that is, 4 bytes for a dot1q tag or 8 bytes for a QinQ SAP). Similarly when a packet is received on a SDP Binding (a.k.a. PW) of type vc-vlan, the service MTU check includes the length of the encapsulated packet along with the vc-vlan encapsulation length.

If the required payload is larger than the port or channel MTU, the SAP transitions to an inoperative state.

If the required MTU is equal to or less than the port or channel MTU, the SAP transitions to an operative state.

The service MTU is compared to the path MTU associated with an SDP before binding an SDP to a service. The path MTU can be administratively defined in the context of the SDP. The default or administrative path MTU can be dynamically reduced based on:

- The MTU capabilities discovered by the tunneling mechanism of the SDP.
- The egress interface MTU capabilities based on the next hop in the tunnel path.

If the service MTU is greater than the path MTU, the SDP binding for the service transitions to an inoperative state. If the service MTU is equal to or less than the path MTU, the SDP binding transitions to an operative state.

If a service MTU, path MTU or a channel MTU is dynamically or administratively modified, the operational states of all associated SAP and SDP bindings are automatically re-evaluated.

The **no** form of the command restores the default service-mtu of the indicated service type to default value.

Notes:

- To disable service MTU check, execute the command **no service-mtu-check**. Disabling service MTU check allows the packets to pass to the egress if the packet length is lesser than or equal to the MTU configured on the port. In 7210 SAS, length of the SAP tag (or service-delimiting tag, for a packet received over a pseudowire) is included in the computation of the packet length before comparing it with the service-MTU configured for the service.
  Packet length= Length of IP packet + L2 header + length of SAP tag

- This command is supported only on 7210 SAS devices configured in Network mode. It is not supported in access-uplink mode.

Default epipe: 1514

The following table displays MTU values for specific VC types.

SAP 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 (Q-in-Q with preserved bottom Qtag)	1518	1504

octets — The size of the MTU in octets, expressed as a decimal integer, between 1 — 9194.

service-name

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

Context config>service>epipe
config>service>cpipe

Description This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a given service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the 7750 SR, 7450 ESS and 7710 SR platforms.

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

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

service-mtu-check

Syntax [**no**] **service-mtu-check**

Context config>service>epipe

Description The **no** form of this command disables the service MTU check.

Disabling service MTU check allows the packets to pass to the egress if the packet length is lesser than or equal to the MTU configured on the port. The length of the packet sent from a SAP is limited only by the access port MTU. In case of a pseudowire the length of a packet is limited by the network port MTU (including the MPLS encapsulation).

Notes:

- If TLDP is used for signaling, the configured value for service-mtu is used during a pseudowire setup.
- This command is supported on 7210 SAS-M and 7210 SAS-T in Network mode.

Default enabled

VLL SAP Commands

sap

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

Context config>service>epipe

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the 7210 device. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP 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.

In a single physical port only one SAP can belong to one service. Multiple SAPs can be defined over a physical port but each of these SAPs should belong to different service.

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.

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:

- Ethernet SAPs support null, dot1q

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 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).

sap-id — Specifies the physical port identifier portion of the SAP. See Common CLI Command Descriptions on page 987 for command syntax.

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

The SAP aggregate rate can be used only if SAP based scheduling mode is configured at the port level. It is not supported in FC-based scheduling mode.

When configured in SAP-based scheduling mode, the egress port scheduler distributes the available bandwidth to all the SAPs configured on the port, up to the configured aggregate rate for the SAP.

tod-suite

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

Context config>service>epipe>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.

accounting-policy

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

Context config>service>epipe>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 1-99

description

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

Context config>service>epipe>sap
config>service>epipe>spoke-sdp

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.

collect-stats

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

Context config>service>cpipe>sap
config>service>cpipe>spoke-sdp
config>service>epipe>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 cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default no collect-stats

ethernet

Syntax **ethernet**

Context config>service>epipe>sap

Description Use this command to configure Ethernet properties in this SAP.

llf

Syntax [**no**] **llf**

Context config>service>epipe>sap>ethernet

Description This command enables Link Loss Forwarding (LLF) on an Ethernet port. It provides an end-to-end OAM fault notification for Ethernet VLL service.LLF on an Ethernet port brings down the port when

there is a local fault on the pseudowire or service, or a remote fault on the SAP or pseudowire, signaled with label withdrawal or TLDP status bits. It ceases when the fault disappears.

The Ethernet port must be configured for null encapsulation.

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

ignore-oper-down

Syntax **[no] ignore-oper-down**

Context config>service>epipe>sap

Description This command enables the user to configure the optional command for a specific SAP to ignore the transition of the operational state to down when a SAP fails. Only a single SAP in an ePipe may use this option.

Default no ignore-oper-down

bit-error-threshold

Syntax **bit-error-threshold** *errors*
no bit-error-threshold

Context config>service>epipe>sap>eth-cfm>mep>eth-test-enable

Description This command is used to specify the threshold value of bit errors.

one-way-delay-threshold

Syntax **one-way-delay-threshold** *seconds*

Context config>service>vpls>sap>eth-cfm>mep

Description This command enables/disables eth-test functionality on MEP.

Parameters *seconds* — Specifies the one way delay threshold in seconds.

Values 0-600

Default 3

mip

Syntax **mip [mac mac-address]**
mip default-mac
no mip

Context config>service>epipe>sap>eth-cfm
config>service>epipe>spoke-sdp>eth-cfm

Description This command allows Maintenance Intermediate Points (MIPs) to be created if mhf-creation for the MA is configured using the default option.

Note: This command is supported on 7210 SAS-X, and 7210 SAS-M, 7210 SAS-T(Network and Access-uplink mode).

Parameters *mac-address* — Specifies the MAC address of the MIP.

Values 6-byte mac-address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MIP. The MAC must be unicast. Using the all zeros address is equivalent to the no form of this command.

default-mac — Using the no command deletes the MIP. If the operator wants to change the mac back to the default mac without having to delete the MIP and reconfiguring this command is useful.

Default no mip

Connection Profile Commands

connection-profile

Syntax **connection-profile** *conn-prof-id [create]*
no connection-profile *conn-prof-id*

Context config

Description This command creates a list of VLAN values to be assigned to a Dot1q SAP in an Epipe service.

A connection profile can only be assigned to a Dot1q SAP which is part of an Epipe Service.

The no form of this command deletes the profile from the configuration.

Default none

Parameters *conn-prof-id* — Specifies the profile number.

Values 1 — 8000

ethernet

Syntax ethernet

Context config>connprof

Description Provides the context to configure the VLAN ranges values.

Default none

ranges

Syntax no ranges
ranges vlan-ranges *[vlan-ranges...(upto 32 max)]*

Context config>connprof>ethernet

Description Specifies the list of VLAN ranges or individual VLAN ID to be used for mapping the given VLANs to the Epipe SAP.

The system validates that the values specified are valid VLAN ID in the range 0-4094 (VLAN ID 4095 is reserved). Ranges are specified in the format ‘a-b ’, the expression (a < b) should be true. Up to about 32 individual VLAN values or VLAN ranges can be specified. A maximum of up to 8 VLAN ranges are allowed per connection profile.

Default none

Parameters *vlan-ranges* — Specifies the list of VLAN ranges or individual VLAN ID to be used for mapping the given VLANs to the Epipe SAP.

Values A list of space separated values specified as either a-b or individual VLAN IDs. Both the VLAN IDs and the value used for 'a' and 'b' must be in the range of 0-4094. Additionally, value 'a' must be less than value 'b'.
For example:

ranges	100-200 5 6 4000-4020
ranges	4 5 6 10 11 12
ranges	250-350 500-600 1000-1023

Service Filter and QoS Policy Commands

egress

Syntax **egress**

Context config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>epipe>sap

Description This command enables the context to configure egress SAP parameters.

force-vlan-vc-forwarding

Syntax [**no**] **force-vlan-vc-forwarding**

Context config>service>epipe>spoke-sdp
config>service>vpls>spoke-sdp

Description This command forces vc-vlan-type forwarding in the data path for spoke which have either vc-type. This comand is not allowed on vlan-vc-type SDPs.

The **no** version of this command sets default behavior.

Default Per default this feature is disabled

ingress

Syntax **ingress**

Context config>service>epipe>sap
config>service>cpipe>sap
config>service>cpipe>spoke-sdp
config>service>epipe>sap>statistics

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.

aggregate-meter-rate

Syntax **aggregate-meter-rate** *rate-in-kbps* [**burst** *burst-in-kbits*]
no aggregate-meter-rate

Context config>service>vpls>sap>ingress

config>service>epipe>sap>ingress

Description This command allows the user to configure the SAP aggregate policer. The rate of the SAP aggregate policer must be specified by the user. The user can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the traffic on different FCs and determines the destination of the packet. The packet is either forwarded to an identified profile or dropped.

Note: The sum of CIR of the individual FCs configured under the SAP cannot exceed the PIR rate configured for the SAP. Though the 7210 SAS software does not block this configuration, it is not recommended for use.

The table below provides information about the final disposition of the packet based on the operating rate of the per FC policer and the per SAP aggregate policer:

Per FC meter Operating Rate	Per FC Assigned Color	SAP aggregate meter Operating Rate	SAP aggregate meter color	Final Packet Color
Within CIR	Green	Within PIR	Green	Green or In-profile
Within CIR*	Green	Above PIR	Red	Green or In-profile
Above CIR, Within PIR	Yellow	Within PIR	Green	Yellow or Out-of-Profile
Above CIR, Within PIR	Yellow	Above PIR	Red	Red or Dropped
Above PIR	Red	Within PIR	Green	Red or Dropped
Above PIR	Red	Above PIR	Red	Red or Dropped

Table 17: Final Disposition of the packet based on per FC and per SAP policer or meter.

Note*: The row number 2 in the above table is not recommended for use. For more information on this, see the Note in the "**aggregate-meter-rate**" description.

When the SAP aggregate policer is configured, per FC policer can be only configured in "trtcm2" mode (RFC 4115).

Note: The meter modes "srtcm" and "trtcm1" are used in the absence of an aggregate meter.

The SAP ingress meter counters increment the packet or octet counts based on the final disposition of the packet.

If ingress Frame-based accounting is used, the SAP aggregate meter rate accounts for the Ethernet frame overhead. The system accounts for 12 bytes of IFG and 8 bytes of start delimiter.

The **no** form of the command removes the aggregate policer from use.

Default no aggregate-meter-rate

Parameters *rate-in-kbps* — Specifies the rate in kilobits per second.

Values 01 — 20000000 | max

Default max

burst <*burst-in-kilobits*> — Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 —2146959

Default 512

filter

Syntax **filter** [**ip** *ip-filter-id*]
filter [**ipv6** *ipv6-filter-id*]
filter [**mac** *mac-filter-id*]
no filter [**ip** *ip-filter-id*]
no filter [**ipv6** *ipv6-filter-id*]
no filter [**mac** *mac-filter-id*]

Context config>service>epipe>sap>egress
config>service>epipe>sap>ingress

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.

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

meter-override

Syntax **[no] meter-override**

Context config>service>epipe>sap>ingress
config>service>vpls>sap>ingress
config>service>ies>interface>sap>ingress
config>service>vprn>interface>sap>ingress

Description This command, within the SAP ingress contexts, is used to create a CLI node for specific overrides to one or more meters created on the SAP through the sap-ingress QoS policies.

The no form of the command is used to remove any existing meter overrides.

Default no meter-overrides

meter

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

Context config>service>epipe>sap>ingress>meter-override
config>service>vpls>sap>ingress>meter-override
config>service>ies>interface>sap>ingress>meter-override
config>service>vprn>interface>sap>ingress>meter-override

Description This command, within the SAP ingress contexts, is used to create a CLI node for specific overrides to a specific meter created on the SAP through a sap-ingress QoS policies.

The no form of the command is used to remove any existing overrides for the specified meter-id.

Parameters *meter-id* — The meter-id parameter is required when executing the meter command within the meter-overrides context. The specified meter-id must exist within the sap-ingress QoS policy applied to the SAP. If the meter is not currently used by any forwarding class or forwarding type mappings, the meter will not actually exist on the SAP. This does not preclude creating an override context for the meter-id.

create — The create keyword is required when a meter *meter-id* override node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

adaptation-rule

Syntax **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]
no adaptation-rule

Context config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter
config>service>vprn>interface>sap>ingress>meter-override>meter

Description This command can be used to override specific attributes of the specified meter adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the meter is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of 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 meter-override meter *meter-id* command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **meter-override** command is not specified, the default applies.

NOTE: When the meter mode in use is 'trtcm2', this parameter is interpreted as EIR value. For more information, see the description and relevant notes for meter modes in the 7210 SAS QoS user guide.

cir — The **cir** parameter defines the constraints enforced when adapting the CIR rate defined within the meter-override meter *meter-id* command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule — Specifies the criteria to use to compute the operational CIR and PIR values for this meter, while maintaining a minimum offset.

Values **max** — The **max** (maximum) keyword is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the meter will be equal to or less than the administrative rate specified using the **meter-override** command.

min — The **min** (minimum) keyword is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **meter-override** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the meter will be the rate closest to the rate specified using the **meter-override** command.

cbs

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

Context config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter
config>service>vprn>interface>sap>ingress>meter-override>meter

Description This command, within the SAP ingress meter-overrides contexts, is used to override the sap-ingress QoS policy configured cbs parameter for the specified meter-id.

The no form of the command is used to restore the meter cbs setting to the meter defined value.

Default no mbs

Parameters *size-in-kbytes* — The size parameter is an integer expression of the number of kilobytes reserved for the meter. 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 [4..2146959 | default]

mbs

Syntax **mbs** *size-in-kbits*
no mbs

Context config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter
config>service>vprn>interface>sap>ingress>meter-override>meter

Description This command, within the SAP ingress meter-overrides contexts, is used to override the sap-ingress QoS policy configured mbs parameter for the specified meter-id.

The no form of the command is used to restore the meter mbs setting to the meter defined value.

Default no mbs

Parameters *size-in-kbits* — The size parameter is required when specifying mbs override and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional byte and kilobyte keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes. When byte is defined, the value given for size is interpreted as the meter MBS value given in bytes. When kilobytes is defined, the value is interpreted as the meter MBS value given in kilobytes.

Values [4..2146959 | default]

mode

Syntax **mode** *mode*
no mode

Context config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter
config>service>vprn>interface>sap>ingress>meter-override>meter

Description This command within the SAP ingress meter-overrides contexts is used to override the sap-ingress QoS policy configured mode parameters for the specified meter-id.

The no mode command is used to restore the policy defined metering and profiling mode to a meter.

Parameters *mode* — Specifies the rate mode of the meter-override.

Values trtcm1|trtcm2|srtcm

rate

Syntax **rate cir** *cir-rate* **[pir** *pir-rate***]**
no rate

Context config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter
config>service>vprn>interface>sap>ingress>meter-override>meter

Description This command within the SAP ingress meter-overrides contexts is used to override the sap-ingress QoS policy configured rate parameters for the specified meter-id.

The no rate command is used to restore the policy defined metering and profiling rate to a meter.

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

Parameters *pir-rate* — 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.

NOTE: When the meter mode is set to 'trtcm2' the PIR value is interpreted as the EIR value. For more information, see the 7210 SAS QoS user guide.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values [0..20000000 | max]

Default max

cir-rate — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been

executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be given as a positive integer.

Values **[0..20000000 | max]**

Default 0

qos

Syntax **qos** *policy-id*
no qos

Context config>service>cpipe>sap>ingress
config>service>epipe>sap>ingress

Description This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP).

QoS ingress 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 ingress . The **qos** command only allows ingress policies to be associated on SAP ingress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type will return an error.

By default, if no specific QoS policy is associated with the SAP for ingress, so the default QoS policy is used.

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

policy-id — The ingress policy ID to associate with SAP on ingress. The policy ID must already exist.

Values 1 — 65535

statistics

Syntax **statistics**

Context config>service>epipe>sap
config>service>vpls>sap

Description This command enables the context to configure the counters associated with SAP ingress and egress.

ingress

Syntax **ingress**

Context config>service>epipe>sap>statistics

config>service>vpls>sap>statistics

Description This command enables the context to configure the ingress SAP statistics counter.

counter-mode

Syntax **counter-mode** {in-out-profile-count| forward-drop-count}

Context

config>service>epipe>sap>statistics>ingress
config>service>vpls>sap>statistics>ingress

Description

This command allows the user to set the counter mode for the counters associated with sap ingress meters (a.ka. policers). A pair of counters is available with each meter. These counters count different events based on the counter mode value.

Note: The counter mode can be changed if an accounting policy is associated with a SAP. If the counter mode is changed the counters associated with the meter are reset and the counts are cleared. If an accounting policy is in use when the counter-mode is changed a new record will be written into the current accounting file.

Note: The configuration information is not saved across re-boot.

Execute the following sequence of commands to ensure a new accounting file is generated when the counter-mode is changed:

1. Execute the command **config>service>epipe/vpls>sap> no collect-stats**, to disable writing of accounting records.
2. Change the counter-mode to the desired value, execute the command **config>service>epipe/vpls>sap>counter-mode {in-out-profile-count| forward-drop-count}**.
3. Execute the command **config>service>epipe/vpls>sap> collect-stats**, to enable writing of accounting records.

The **no** form of the command restores the counter mode to the default value.

Default when either in-out-profile-count or forward-drop-count is in use in-out-profile-count

Parameters **forward-drop-count —** If the counter mode is specified as "forward-drop-count", one counter counts the forwarded packets and octets received on ingress of a SAP and another counts the dropped packets. The forwarded count is the sum of in-profile and out-of-profile packets/octets received on SAP ingress. The dropped count is count of packets/octets dropped by the policer. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. The in-profile count and out-of-profile count is not individually available when operating in this mode.

in-out-profile-count — If the counter mode is specified as "in-out-profile-count", one counter counts the total in-profile packets and octets received on ingress of a SAP and another counts the total out-of-profile packets and octets received on ingress of a SAP. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. Dropped counts are not maintained in hardware when this mode is used. It is obtained by subtracting the sum of in-profile count and out-of-profile count from the total SAP ingress received count and displayed.

VLL SDP Commands

Note : VLL SDP commands are not supported on 7210 SAS-M and 7210 SAS-T devices configured in access uplink mode.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**no-endpoint**] [**create**]
spoke-sdp *sdp-id*[**:***vc-id*] **endpoint** *endpoint-name*
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>cpipe
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 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 or VPL 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 7210 SAS M 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. Vc-switching VLLs are an exception. If the VLL is a "vc-switching" VLL, then the two endpoints must both be SDPs.

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 — Removes the association of a spoke SDP with an explicit endpoint name.

endpoint *endpoint-name* **—** Specifies the name of the service endpoint.

control-word

Syntax [**no**] **control-word**

Context config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp

Description The control word command provides the option to add a control word as part of the packet encapsulation for pseudowire types for which the control word is optional. These are Ethernet pseudowires (Epipe).

The configuration for the two directions of the pseudowire must match because the control word negotiation procedures described in Section 6.2 of RFC 4447 are not supported. The C-bit in the pseudowire FEC sent in the label mapping message is set to 1 when the control word is enabled. Otherwise, it is set to 0.

The service will only come up if the same C-bit value is signaled in both directions. If a spoke-sdp is configured to use the control word but the node receives a label mapping message with a C-bit clear, the node releases the label with the an “Illegal C-bit” status code as per Section 6.1 of RFC 4447. As soon as the user also enabled the control the remote peer, the remote peer will withdraw its original label and will send a label mapping with the C-bit set to 1 and the VLL service will be up in both nodes.

control-channel-status

Syntax [**no**] **control-channel-status**

Context config>service>epipe>spoke-sdp

Description **Platforms Supported:** 7210 SAS-T and 7210 SAS-R6.

This command enables the configuration of static pseudowire status signaling on a spoke-sdp for which signaling for its SDP is set to OFF.

A control-channel-status no shutdown is allowed only if all of the following is true:

- The system is using network chassis mode D
- SDP signaling is off
- The control-word is enabled (control-word by default is disabled)
- The service type is epipe, apipe, vpls, cpipe, or IES/VPRN
- Mate sdp signaling is off (in vc-switched services)
- pw-path-id is configured for this spoke

The **no** form of this command renoves control channel status signaling form a spoke-sdp. It can only be removed if control channel status is shutdown.

Default no control-channel-status

acknowledgment

Syntax [**no] acknowledgment**

Context config>service>epipe>spoke-sdp>control-channel-status

Description Platforms Supported: 7210 SAS-T and 7210 SAS-R6.

This command enables the acknowledgement of control channel status messages. By default, no acknowledgement packets are sent.

refresh-timer

Syntax **refresh-timer** *value*
no refresh-timer

Context config>service>epipe>spoke-sdp>control-channel-status

Description **Platforms Supported:** 7210 SAS-T and 7210 SAS-R6.

This command configures the refresh timer for control channel status signaling packets. By default, no referesh packets are sent.

Default no refresh-timer

Parameters *value* — Specifies the refresh timer value.

Values 10 — 65535 seconds

Default 0 (off)

request-timer

Syntax **request-timer** request-timer *request-timer-secs* retry-timer *retry-timer-secs* timeout-multiplier *multiplier*

Context config>service>epipe>spoke-sdp>control-channel-status

Description **Platforms Supported**: 7210 SAS-T and 7210 SAS-R6.

This command configures the control channel status request mechanism. When it is configured, control channel status request procedures are used. These augment the procedures for control channel status messaging from RFC 6478. This command is mutually exclusive with a non-zero refresh-timer value.

Parameters *request-timer* — Specifies the interval at which pseudowire status messages, including a reliable delivery TLV, with the “request” bit set, are sent.

Values 10 — 65535 seconds

retry-timer — specifies the timeout interval if no response to a pseudowire status request is received. This parameter must be configured. A value of zero (0) disables retries.

Values 0, 3 — 60 seconds

timeout-multiplier — If a requesting node does not receive a valid response to a pseudowire status request within this multiplier times the retry timer, then it will assume the pseudowire is down. This parameter is optional.

Values 3 — 20 seconds

precedence

Syntax **precedence** [*precedence-value* | **primary**]
no precedence

Context config>service>cpipe>spoke-sdp
config>service>epipe>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.

pw-path-id

Syntax [**no**] **pw-path-id**

Context config>service>epipe>spoke-sdp

Description **Platforms Supported:** 7210 SAS-T and 7210 SAS-R6.

This command enables the context to configure an MPLS-TP Pseudowire Path Identifier for a spoke-sdp. All elements of the PW path ID must be configured in order to enable a spoke-sdp with a PW path ID.

For an IES or VPRN spoke-sdp, the pw-path-id is only valid for ethernet spoke-sdps.

The **pw-path-id** only configurable if all of the following is true:

- The system is using network chassis mode D
- SDP signaling is off
- control-word is enabled (control-word is disabled by default)
- the service type is epipe, vpls, cpipe, or IES/VPRN interface
- mate SDP signaling is off for vc-switched services

The **no** form of the command deletes the PW path ID.

Default no pw-path-id

agi

Syntax **agi** *agi*
no agi

Context config>service>epipe>spoke-sdp>pw-path-id

Description **Platforms Supported:** 7210 SAS-T and 7210 SAS-R6.

This command configures the attachment group identifier for an MPLS-TP PW.

Parameters *agi* — Specifies the attachment group identifier.

Values 0 — 4294967295

saii-type2

Syntax **saii-type2** *global-id:node-id:ac-id*
no saii-type2

Context config>service>epipe>spoke-sdp>pw-path-id

Description **Platforms Supported:** 7210 SAS-T and 7210 SAS-R6.

This command configures the source individual attachement identifier (SAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured, that is, if it is at an S-PE, then the values must match those of the taii-type2 of the mate spoke-sdp.

Parameters *global-id* — Specifies the global ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values 0 — 4294967295

node-id — Specifies the node ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values a.b.c.d or 0 — 4294967295

ac-id — Specifies the attachment circuit ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value.

Values 1 — 4294967295

taii-type2

Syntax **taii-type2** *global-id:node-id:ac-id*
no taii-type2

Context config>service>epipe>spoke-sdp>pw-path-id

Description **Platforms Supported:** 7210 SAS-T and 7210 SAS-R6.

This command configures the source individual attachement identifier (SAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured i.e. it is at an S-PE, then the values must match those of the taii-type2 of the mate spoke-sdp.

Parameters *global-id* — Specifies the global ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values 0 — 4294967295

node-id — Specifies the node ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values a.b.c.d or 0 — 4294967295

ac-id — Specifies the attachment circuit ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value.

Values 1 — 4294967295

pw-status-signaling

Syntax [**no**] **pw-status-signaling**

Context config>service>epipe>spoke-sdp

Description **Platforms Supported:** 7210 SAS-T and 7210 SAS-R6.

This command enables pseudowire status signaling for this spoke SDP binding.

The **no** form of the command disables the status signaling.

Default pw-status-signaling

vc-label

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

Context config>service>cpipe>spoke-sdp>egress
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>cpipe>spoke-sdp>ingress
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.

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) 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. The 7210 SAS supports provisioning of access or uplink spokes to connect to the provider edge IP/MPLS routers.

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) service routers. The connection appears to be a bridged domain to the customer sites so protocols, including routing protocols, can traverse the VPLS service.

Other VPLS advantages include:

- VPLS is a transparent, protocol-independent service.
- There is no Layer 2 protocol conversion between LAN and WAN technologies.
- There is no need to design, manage, configure, and maintain separate WAN access equipment, thus, eliminating the need to train personnel on WAN technologies such as Frame Relay.

VPLS Packet Walkthrough in Network Mode

This section provides an example of VPLS processing of a customer packet sent across the network from site-A, which is connected to PE-Router-A through a 7210 SAS to site-C, which is connected through 7210 SAS to PE-Router-C (Figure 33) in an HVPLS configuration. This section does not discuss the processing on the PE routers, but only on 7210 SAS routers.

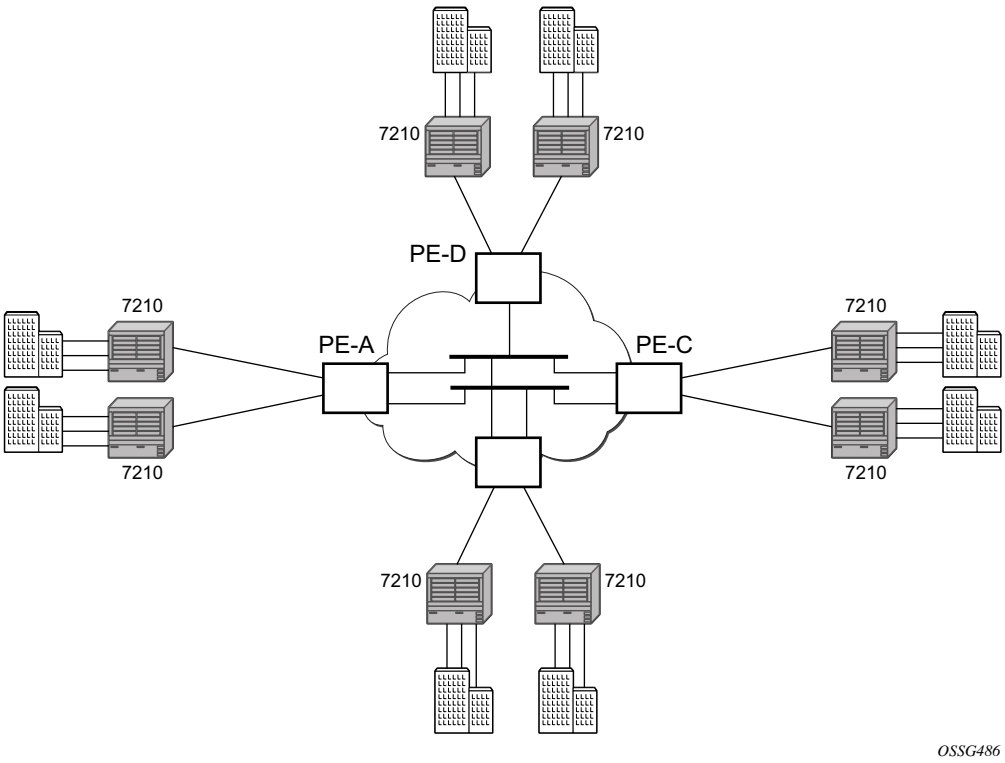

Figure 33: VPLS Service Architecture

1. 7210-A (Figure 34)
    a. Service packets arriving at 7210-A are associated with a VPLS service instance based on the combination of the physical port and the IEEE 802.1Q tag (VLAN-ID) in the packet

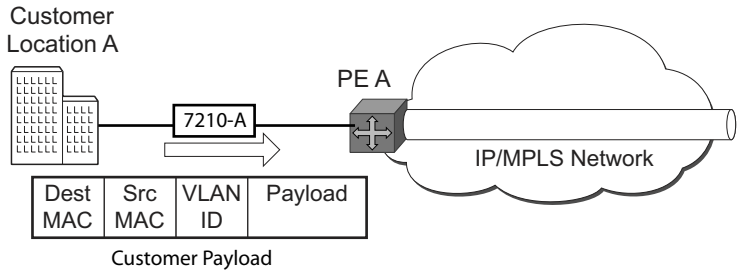

Figure 34: Access Port Ingress Packet Format and Lookup

b. 7210-A learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address to the service access point (SAP) on which it was received.

c. The destination MAC address in the packet is looked up in the FIB table for the VPLS instance. There are two possibilities: either the destination MAC address has already been learned (known MAC address) or the destination MAC address is not yet learned (unknown MAC address).

   For a Known MAC Address (Figure 35):

d. If the destination MAC address has already been learned by 7210, an existing entry in the FIB table identifies the far-end PE-Router and the service VC-label (inner label) to be used before sending the packet to PE-Router-A.

e. The customer packet is sent on this LSP 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 35):

f. If the destination MAC address has not been learned, 7210 will flood the packet to spoke SDPs that are participating in the service.

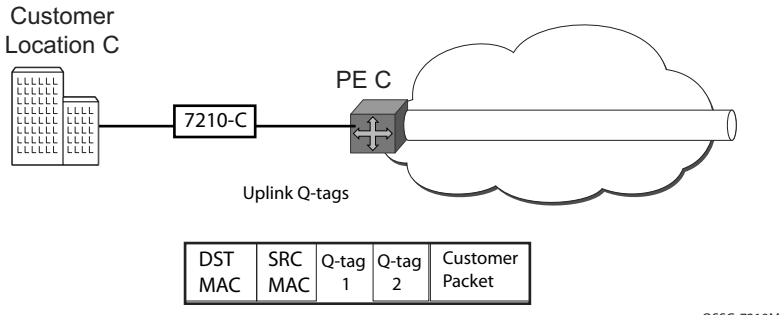

Figure 35: Network Port Egress Packet Format and Flooding

2. Core Router Switching

a. The PE router will encapsulate this packet in the appropriate MPLS header and transport it across the core network to the remote 7210-C.

3. 7210-C (Figure 34)

   a. 7210-C associates the packet with the VPLS instance based on the VC label in the received packet after the stripping of the tunnel label.

   b. 7210-C learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address to the spoke SDP on which the packet was received.

   c. The destination MAC address in the packet is looked up in the FIB table for the VPLS instance. Again, there are two possibilities: either the destination MAC address has already been learned (known MAC address) or the destination MAC address has not been learned on the access side of 7210-C (unknown MAC address).

   d. If the destination MAC address has been learned by 7210-C, an existing entry in the FIB table identifies the local access port and the IEEE 802.1Q tag (if any) to be added before sending the packet to customer Location-C. The egress Q tag may be different than the ingress Q tag.

   e. If the destination MAC address has not been learned, 7210 will flood the packet to all the access SAPs that are participating in the service.

VPLS Packet Walkthrough in Access Uplink Mode

This section provides an example of VPLS processing of a customer packet sent across the network from site-A, which is connected to PE-Router-A through a 7210 SAS M to site-C, which is connected through 7210 SAS M to PE-Router-C (Figure 33) in an HVPLS configuration. This section does not discuss the processing on the PE routers, but only on 7210 SAS routers.

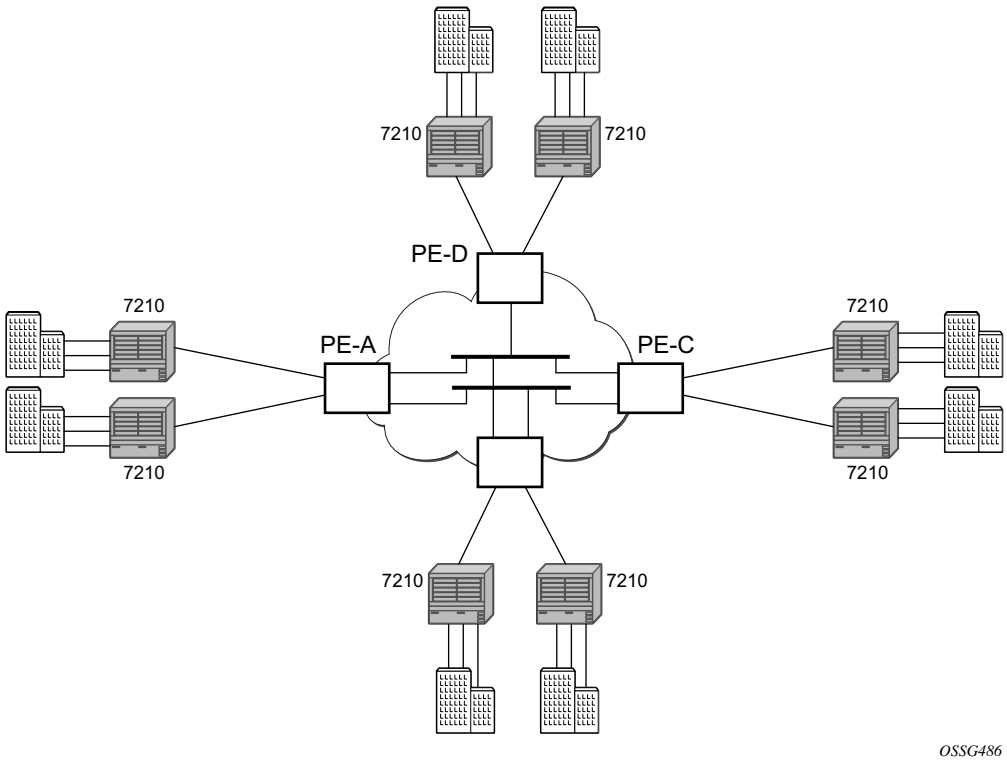

Figure 36: VPLS Service Architecture

1. 7210-A (Figure 34)
    a. Service packets arriving at 7210-A are associated with a VPLS service instance based on the combination of the physical port and the IEEE 802.1Q tag (VLAN-ID) in the packet.

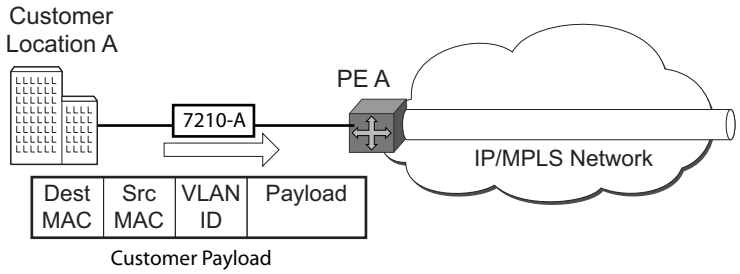

Figure 37: Access Port Ingress Packet Format and Lookup

b. 7210-A learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address to the service access point (SAP) on which it was received.

c. The destination MAC address in the packet is looked up in the FIB table for the VPLS instance. There are two possibilities: either the destination MAC address has already been learned (known MAC address) or the destination MAC address is not yet learned (unknown MAC address).

For a Known MAC Address (Figure 35):

d. If the destination MAC address has already been learned by 7210, an existing entry in the FIB table identifies destination uplink QinQ SAP to be used for sending the packet towards the PE-Router-A.

e. The customer packet is sent on this uplink SAP once the IEEE 802.1Q tag is stripped and the uplink SAP tag is added to the packet.

For an Unknown MAC Address (Figure 35):

f. If the destination MAC address has not been learned, 7210 will flood the packet to all the uplink SAPsspoke SDPs that are participating in the service .

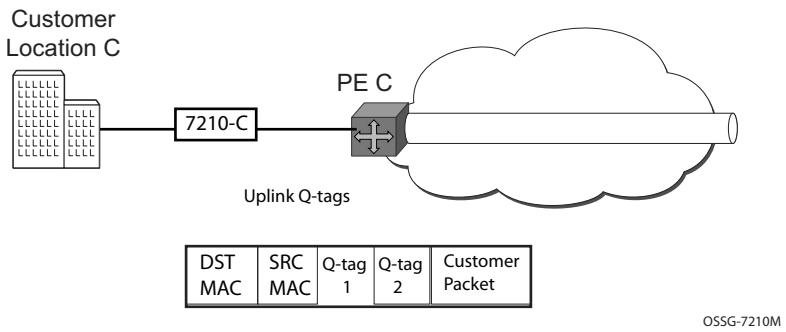

Figure 38: Network Port Egress Packet Format and Flooding

2. Core Router Switching

a. The PE router will encapsulate this packet in the appropriate MPLS header and transport it across the core network to the remote 7210-C.

3. 7210-C (Figure 34)

a. 7210-C associates the packet with the VPLS instance based on the VLAN tags in the received packet.

b. 7210-C learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address to the access uplink porton which the packet was received.

c. The destination MAC address in the packet is looked up in the FIB table for the VPLS instance. Again, there are two possibilities: either the destination MAC address has already been learned (known MAC address) or the destination MAC address has not been learned on the access side of 7210-C (unknown MAC address).

d. If the destination MAC address has been learned by 7210-C, an existing entry in the FIB table identifies the local access port and the IEEE 802.1Q tag (if any) to be added before sending the packet to customer Location-C. The egress Q tag may be different than the ingress Q tag.

e. If the destination MAC address has not been learned, 7210 will flood the packet to all the access SAPs that are participating in the service.

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 including:
  - → Configurable FIB size limit
  - → FIB size alarms
  - → MAC learning disable
  - → Discard unknown
  - → Separate aging timers for locally and remotely learned MAC addresses.
- Ingress rate limiting for broadcast, multicast, and destination unknown flooding on a per SAP basis.
- Implementation of Spanning Tree Protocol (STP) parameters on a per VPLS, per SAP and per spoke SDP basis.
- Optional SAP and/or spoke SDP redundancy to protect against node failure.
- IGMP snooping on a per-SAP and SDP basis.

VPLS over MPLS in Network Mode

The VPLS architecture proposed in *draft-ietf-ppvpn-vpls-ldp-0x.txt* specifies the use of provider equipment (PE) that is capable of learning, bridging, and replication on a per-VPLS basis. The PE routers that participate in the service are connected using MPLS Label Switched Path (LSP) tunnels in a full-mesh composed of mesh SDPs or based on an LSP hierarchy (Hierarchical VPLS (H-VPLS)) composed of mesh SDPs and spoke SDPs. The 7210 SAS M supports only H-VPLS.

Multiple VPLS services can be offered over the same set of LSP tunnels. Signaling specified in *RFC 4905* is used to negotiate a set of ingress and egress VC labels on a per-service basis. The VC labels are used by the PE routers for de-multiplexing traffic arriving from different VPLS services over the same set of LSP tunnels.

VPLS/HVPLS is provided over MPLS by:

- Connecting 7210 SAS M to bridging-capable provider edge (PE) routers through a mesh/spoke SDP. The PE routers are connected using a full mesh of LSPs.
- Negotiating per-service VC labels using draft-Martini encapsulation.
- Replicating unknown and broadcast traffic in a service domain.
- Enabling MAC learning over tunnel and access ports (see VPLS MAC Learning and Packet Forwarding on page 274).
- Using a separate forwarding information base (FIB) per VPLS service.

VPLS over QinQ Spokes for 7210 SAS devices Configured in Access Uplink Mode

7210 SAS devices configured in uplink mode support QinQ spokes or Dot1q spokes, which allows them to connect to upstream PE nodes which provides IP/MPLS transport.

VPLS is provided over QinQ/Dot1q spokes by:

- Connecting bridging-capable 7210 SAS devices.
- Replicating unknown and broadcast traffic in a service domain.
- Enabling MAC learning over QinQ/Dot1q spokes and access ports (see VPLS MAC Learning and Packet Forwarding).
- Using a separate forwarding information base (FIB) per VPLS service.

VPLS MAC Learning and Packet Forwarding

The 7210 SAS edge devices perform the packet replication required for broadcast and multicast traffic across the bridged domain. MAC address learning is performed by the 7210 SAS device to reduce the amount of unknown destination MAC address flooding.

Each 7210 SAS maintains a Forwarding Information Base (FIB) for each VPLS service instance and learned MAC addresses are populated in the FIB table of the service. All traffic is switched based on MAC addresses and forwarded between all participating nodes using the LSP tunnels Unknown destination packets (for example, the destination MAC address has not been learned) are forwarded on all LSPs to all participating nodes for that service until the target station responds and the MAC address is learned by the 7210 SAS associated with that service.

IGMP Snooping in Network Mode and Access-uplink Mode

In Layer 2 switches, multicast traffic is treated like an unknown MAC address or broadcast frame, which causes the incoming frame to be flooded out (broadcast) on every port within a VLAN. Although this is acceptable behavior for unknowns and broadcast frames, this flooded multicast traffic may result in wasted bandwidth on network segments and end stations, as IP multicast hosts can join and be interested in only specific multicast groups.

IGMP snooping entails using information in Layer 3 protocol headers of multicast control messages to determine the processing at Layer 2. By doing so, an IGMP snooping switch provides the benefit of conserving bandwidth on those segments of the network in which no node has expressed interest in receiving packets addressed to the group address.

Note: In the following paragraph on IGMP snooping, reference to SDP is applicable only in network mode.

IGMP snooping can be enabled in the context of VPLS services. The IGMP snooping feature allows for optimization of the multicast data flow to only those SAPs or SDPs that are members of the group. The system builds a database of group members per service by listening to IGMP queries and reports from each SAP or SDP:

- When the switch receives an IGMP report from a host for a particular multicast group, the switch adds the host port number to the forwarding table entry.
- When it receives an IGMP leave message from a host, it removes the host port from the table entry, if no other group members are present. It also deletes entries if it does not receive periodic IGMP membership reports from the multicast clients.

The following are IGMP snooping features:

- IGMP v1, v2, and v3 are supported (RFC 1112, *Host Extensions for IP Multicasting*, and RFC 2236, *Internet Group Management Protocol, Version 2*). 7210 SAS-M and 7210 SAS-T supports IGMPv3 in access-uplink mode. 7210 SAS-M and 7210 SAS-T in network mode does not support IGMPv3.
- IGMP snooping can be enabled and disabled on individual VPLS service instances.
- IGMP snooping can be configured on individual SAPs that are part of a VPLS service. When IGMP snooping is enabled on a VPLS service, all its contained SAPs and SDPs automatically have snooping enabled.
- Fast leave terminates the multicast session immediately, rather than using the standard group-specific query to check if other group members are present on the network.
- SAPs and SDPs can be statically configured as multicast router ports. This allows the operator to control the set of ports to which IGMP membership reports are forwarded.
- Static multicast group membership on a per SAP and as per SDP basis can be configured.

- The maximum number of multicast groups (static and dynamic) that a SAP or SDP can join can be configured. An event is generated when the limit is reached.
- The maximum number of multicast groups (static and dynamic) that a VPLS instance simultaneously supports can be configured.
- Proxy summarization of IGMP messages reduces the number of IGMP messages processed by upstream devices in the network.
- IGMP filtering allows a subscriber to a service or the provider to block, receive, or transmit permission (or both) to individual hosts or a range of hosts.
  The following types of filters can be defined:
  - → Filter group membership that report from a particular host or range of hosts. This filtering is performed by importing an appropriately-defined routing policy into the SAP or SDP.
  - → Filters that prevent a host from transmitting multicast streams into the network. The operator can define a data-plane filter (ACL) that drops all multicast traffic, and apply this filter to a SAP or SDP.

Configuration Guidelines for IGMP Snooping

The following IGMP snooping considerations apply:

- Layer 2 multicast is supported in VPLS services.
- IGMP snooping is not supported for VCs (either vc-ether or vc-vlan) with control-word enabled.
- IGMP snooping fast leave processing can be enabled only on SAPs and SDPs. IGMP snooping proxy summarization is enabled by default on SAPS and SDPs and cannot be disabled. Proxy summarization and fast leave processing are supported only on SDPs whose VC are configured to use vc-type ether and do not have control-word enabled.
- IGMP filtering using policies is available on SAPs and SDPs. It is supported only on SDPs whose VC are configured to use vc-type ether and do not have control-word enabled.
- Dynamic learning is only supported on SDPs whose VC are configured to use vc-type ether and do not have control-word enabled.
- SDPs that are configured to use VC of type 'vc-vlan' that need to be mrouter ports must be configured statically. Multicast group memberships for such SDPs must be configued statically. Dynamic learning is not available for these SDPs.
- IGMP snooping is not supported for control word enabled SDP.
- 7210 SAS-M and 7210 SAS-T in network mode does not support IGMPv3.

Multicast VLAN Registration (MVR) support

Multicast VPLS Registration (MVR) is a bandwidth optimization method for multicast in a broadband services network. MVR allows a subscriber on a port to subscribe and unsubscribe to a multicast stream on one or more network-wide multicast VPLS instances.

MVR assumes that subscribers join and leave multicast streams by sending IGMP join and leave messages. The IGMP leave and join message are sent inside the VPLS to which the subscriber port is assigned. The multicast VPLS is shared in the network while the subscribers remain in separate VPLS services. Using MVR, users on different VPLS cannot exchange any information between them, but still multicast services are provided.

On the MVR VPLS, IGMP snooping must be enabled. On the user VPLS, IGMP snooping and MVR work independently. If IGMP snooping and MVR are both enabled, MVR reacts only to join and leave messages from multicast groups configured under MVR. Join and leave messages from all other multicast groups are managed by IGMP snooping in the local VPLS. This way, potentially several MVR VPLS instances could be configured, each with its own set of multicast channels.

MVR by proxy — In some situations, the multicast traffic should not be copied from the MVR VPLS to the SAP on which the IGMP message was received (standard MVR behavior) but to another SAP. This is called MVR by proxy.

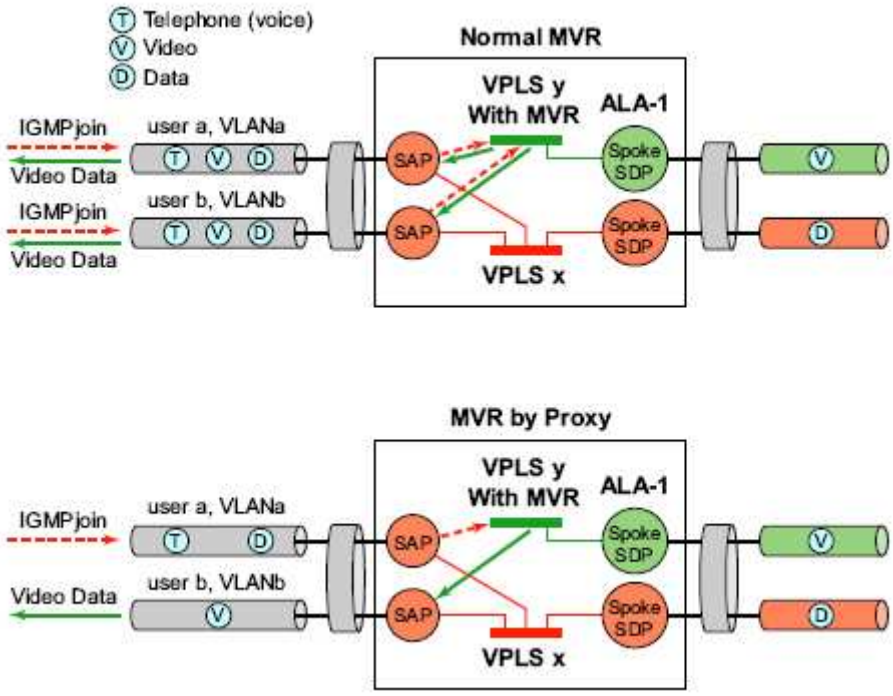

Figure 39: MVR and MVR by Proxy

Configuration Guidelines for MVR

In a MVR configuration, the svc-sap-type of the VPLS service that is the source, which is also known as 'mvr vpls service' and the svc-sap-type of the VPLS service that is the sink, which is also known as 'user vpls service' should match.

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 remotely for a spoke SDP. If the configured limit is reached, then no new addresses will be learned from the SAP or spoke SDP until at least one FIB entry is aged out or cleared.
  - → When the limit is reached on a SAP or spoke SDP, packets with unknown source MAC addresses are still forwarded (this default behavior can be changed by configuration). By default, if the destination MAC address is known, it is forwarded based on the FIB, and if the destination MAC address is unknown, it will be flooded. Alternatively, if discard unknown is enabled at the VPLS service level, unknown destination MAC addresses are discarded.
  - → The log event SAP MAC limit reached is generated when the limit is reached. When the condition is cleared, the log event SAP MAC Limit Reached Condition Cleared is generated.
  - → Disable learning at the VPLS service level allows users to disable the dynamic learning function on the service. Disable Learning is supported at the SAP and spoke SDP level as well.
  - → Disable aging allows users to turn off aging for learned MAC addresses. It is supported at the VPLS service level, SAP level and spoke SDP level

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 by an SDP from another router for the VPLS instance. The local-age timer for the VPLS instance specifies the aging time for locally learned MAC addresses, and the remote-age timer specifies the aging time for remotely learned MAC addresses.

In general, the remote-age timer is set to a longer period than the local-age timer to reduce the amount of flooding required for destination unknown MAC addresses. The aging mechanism is considered a low priority process. In most situations, the aging out of MAC addresses can happen in within tens of seconds beyond the age time. To minimize overhead, local MAC addresses on a LAG port and remote MAC addresses, in some circumstances, can take up to two times their respective age timer to be aged out.

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 for a service, new source MAC addresses are not entered in the VPLS FIB. MAC learning can be disabled for individual SAPs or spoke SDPs.

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 meters can be defined per forwarding class to provide rate-limiting/policing 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 7210 SAS M will shut down the SAP or spoke SDP and create an alarm event when the threshold is exceeded.

MAC move allows sequential order port blocking. By configuration, some VPLS ports can be configured as “non-blockable” which allows simple level of control which ports are being blocked during loop occurrence.

Split Horizon SAP Groups and Split Horizon Spoke SDP Groups

Note: Split Horizon group per service is supported only on 7210 SAS-M and 7210 SAS-T devices configured in Network mode.

Within the context of VPLS services, a loop-free topology inside a fully meshed VPLS core is achieved by applying a split-horizon forwarding concept.The packets received from a mesh SDP are never forwarded to other mesh SDPs within the same service. The advantage of this approach is that no protocol is required to detect loops within the VPLS core network.

In applications such as DSL aggregation, it is useful to extend this split-horizon concept also to groups of SAPs and/or spoke SDPs. This extension is referred to as a split horizon SAP group . Traffic arriving on a SAP or a spoke SDP within a split horizon group will not be forwarded to other SAPs and spoke SDPs configured in the same split horizon group, but will be forwarded to other SAPs/spoke SDPs, which are not part of the split horizon group.

Configuration Guidelines for use of Split Horizon Group in a VPLS Service

In 7210 SAS devices, mesh SDPs cannot be configured in a service which uses split horizon group. Conversely, if a service has a mesh-sdp configured, split horizon group cannot be used in the same service.

Only one split horizon group per service is allowed for use.

VPLS and Spanning Tree Protocol

Alcatel-Lucent's VPLS service provides a bridged or switched Ethernet Layer 2 network. Equipment connected to SAPs forward Ethernet packets into the VPLS service. The 7210 SAS participating in the service learns where the customer MAC addresses reside, on ingress SAPs.

Unknown destinations, broadcasts, and multicasts are flooded to all other SAPs in the service. If SAPs are connected together, either through misconfiguration or for redundancy purposes, loops can form and flooded packets can keep flowing through the network. 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 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 329.

Spanning Tree Operating Modes

Per VPLS instance, a preferred STP variant can be configured. The STP variants supported are:

- `rstp` — Rapid Spanning Tree Protocol (RSTP) compliant with IEEE 802.1D-2004 - default mode
- `dot1w` — Compliant with IEEE 802.1w
- `comp-dot1w` — Operation as in RSTP but backwards compatible with IEEE 802.1w (this mode allows interoperability with some MTU types)
- `mstp` — Compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q-REV/D5.0-09/2005. This mode of operation is only supported in an mVPLS.

While the 7210 SAS initially uses the mode configured for the VPLS, it will dynamically fall back (on a per-SAP basis) to STP (IEEE 802.1D-1998) based on the detection of a BPDU of a different format. A trap or log entry is generated for every change in spanning tree variant.

Some older 802.1W compliant RSTP implementations may have problems with some of the features added in the 802.1D-2004 standard. 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 7210 SAS supports two BDPU encapsulation formats, and can dynamically switch between the following supported formats (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 7210 SAS implementation of Management VPLS (mVPLS) is used to group different VPLS instances under single RSTP instance. Introducing MSTP into the mVPLS allows the following:

- Interoperation with traditional Layer 2 switches in access network.
- 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 the following:

- Configuring mVPLS on VPLS-PEs (only PEs directly connected to GigE MAN network).
- Assign different managed-vlan ranges to different MSTP instances.

Typically, the mVPLS would have SAPs with null encapsulations (to receive, send, and transmit MSTP BPDUs) and a mesh SDP to interconnect a pair of VPLS PEs.

Different access scenarios are displayed in Figure 40 as example network diagrams dually connected to the PBB PEs:

- **Access Type A** — Source devices connected by null or Dot1q SAPs
- **Access Type B** — One QinQ switch connected by QinQ/801ad SAPs
- **Access Type C** — Two or more ES devices connected by QinQ/802.1ad SAPs

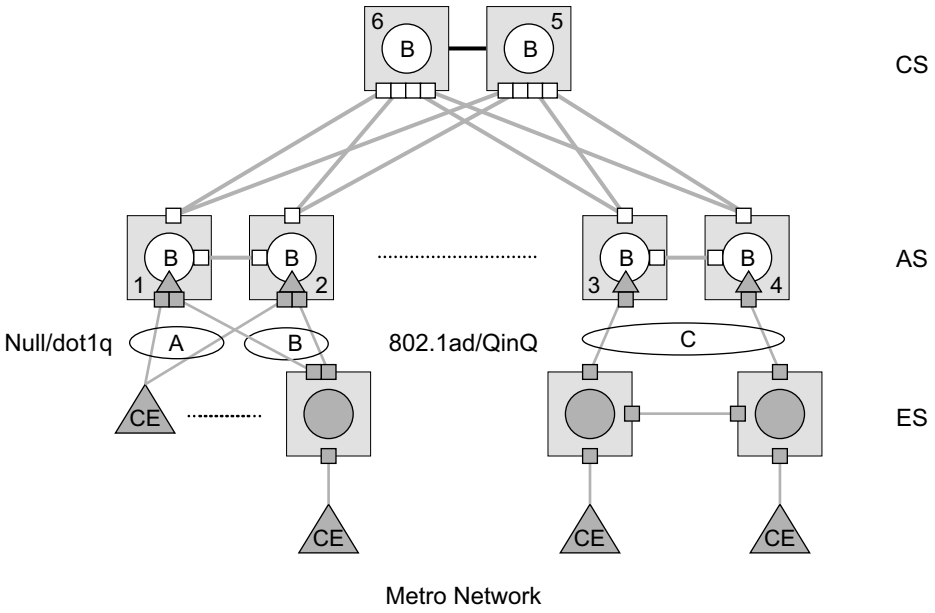

Figure 40: Access Resiliency

The following mechanisms are supported for the I-VPLS:

- **STP/RSTP** can be used for all access types
- **M-VPLS with MSTP** can be used as is just for access Type A. MSTP is required for access type B and C.
- **LAG and MC-LAG** can be used for access Type A and B.
- **Split-horizon-group** does not require residential.

MSTP for QinQ SAPs

MSTP runs in a MVPLS context and can control SAPs from source VPLS instances. QinQ SAPs are supported. The outer tag is considered by MSTP as part of VLAN range control

---

Provider MSTP

Provider MSTP is specified in (IEEE-802.1ad-2005). It uses a provider bridge group address instead of a regular bridge group address used by STP, RSTP, MSTP BPDUs. This allows for implicit separation of source and provider control planes.

The 802.1ad access network sends PBB PE P-MSTP BPDUs using the specified MAC address and also works over QinQ interfaces. P-MSTP mode is used in PBBN for core resiliency and loop avoidance.

Similar to regular MSTP, the STP mode (for example, PMSTP) is only supported in VPLS services where the m-VPLS flag is configured.

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 Layer 2 switches running MSTP protocol assign all VLANs to the IST instance per default. The operator may then “re-assign” individual VLANs to a given MSTI by configuring per VLAN assignment. This means that a SR-Series PE can be considered as the part of the same MST region only if the VLAN assignment to IST and MSTIs is identical to the one of Layer 2 switches in access network.

MSTP in the 7210 SAS Platform

The 7210 SAS platform uses a concept of mVPLS to group different SAPs under a single STP instance. The VLAN range covering SAPs to be managed by a 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 7210 SAS devices (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 7210 SAS devices is closest to the root of the network. This PE device is internally designated as the primary bridge for the VPLS mesh. As a result of this, all network ports on the primary bridges are assigned the designated port role and therefore remain in the forwarding state.

The second part of the solution ensures that the remaining PE devices participating in the STP instance see the SDP ports as a lower cost path to the root 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 are seen as lower cost than any alternative and the PE node will designate the network port as the root port. This ensures that network ports always remain in forwarding state.

A combination of the above mentioned features ensure that network ports are never blocked and maintain interoperability with bridges external to the mesh that are running STP instances.

L2PT Termination

L2PT is used to transparently transport protocol data units (PDUs) of Layer 2 protocols such as STP, CDP, DTP, VTP, PAGP, and UDLD. This allows running these protocols between customer CPEs without involving backbone infrastructure.

The 7210 SAS routers allow transparent tunneling of PDUs across the VPLS core. However, in some network designs, the VPLS PE is connected to CPEs through a legacy Layer 2 network, 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.

The 7210 SAS nodes support L2PT termination for STP BPDUs. 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. The type of protocol can be derived from LLC and SNAP encapsulation.
- In egress direction, PDUs of the corresponding protocol received on all VPLS ports will be intercepted and L2PT encapsulation will be performed for 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) 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 necessary to provide BPDU translation in order to provide an interoperable e2e solution.

To address these network designs, BPDU format translation is supported on 7210 SAS M 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 the VPLS ports will be redirected to the CPM.

BPDU translation involves all encapsulation actions that the data path would perform for a given outgoing port (such as adding VLAN tags depending on the outer SAP and the SDP encapsulation type) and adding or removing all the required VLAN information in a BPDU payload.

This feature can be enabled on a SAP/spoke only if STP is disabled in the context of the given VPLS service.

L2PT and BPDU Translation

Tunnelling of Cisco Discovery Protocol (CDP), Digital Trunking Protocol (DTP), Port Aggregation Protocol(PAGP), Uni-directional Link Detection (UDLD), Virtual Trunk Protocol (VTP), STP (Spanning Tree Protocol) and PVST (per-VLAN Spanning Tree protocol) are supported. These protocols automatically pass the other protocols tunneled by L2PT towards the CPM and on tunnelling, all carry the same specific Cisco MAC.

The existing L2PT limitations apply.

- The protocols apply only to VPLS.
- The protocols are mutually exclusive with running STP on the same VPLS as soon as one SAP/spoke has L2PT/BPDU translation enabled.
- Forwarding occurs on the CPM.

VPLS Redundancy

The VPLS standard (RFC 4762, *Virtual Private LAN Services Using LDP Signalling*) includes provisions for hierarchical VPLS, using point-to-point spoke SDPs. Two applications have been identified for spoke SDPs:

- To connect to Multi-Tenant Units (MTUs) to PEs in a metro area network;
- To interconnect the VPLS nodes of two networks.

In both applications the spoke SDPs serve to improve the scalability of VPLS. While node redundancy is implicit in non-hierarchical VPLS services (using a full mesh of SDPs between PEs), node redundancy for spoke SDPs needs to be provided separately. In VPLS services,only two spoke-SDPs are allowed in an endpoint.

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 41), 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 node 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 VPLS services. The VPLS instance running STP is referred to as the “management VPLS” or mVPLS.

In the case of a failure of the active node, STP on the management VPLS in the standby node will change the link states from disabled to active. The standby node will then broadcast a MAC flush LDP control message in each of the protected VPLS instances, so that the address of the newly active node can be re-learned by all PEs in the VPLS.

It is possible to configure two management VPLS services, where both VPLS services have different active spokes (this is achieved by changing the path-cost in STP). By associating different user VPLSes with the two management VPLS services, load balancing across the spokes can be achieved.

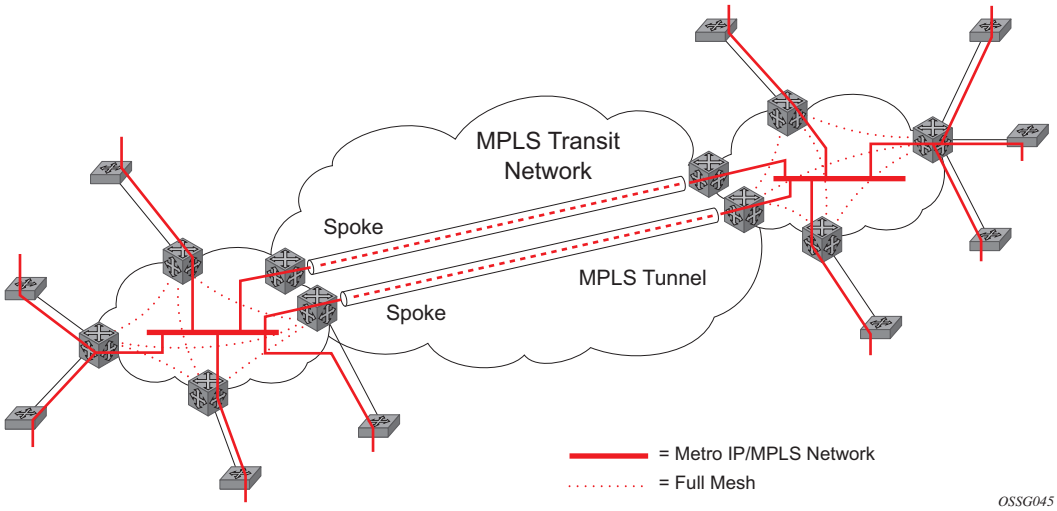

Figure 41: HVPLS with Spoke Redundancy

Spoke SDP Based Redundant Access

This feature provides the ability to have a node deployed as MTUs (Multi-Tenant Unit Switches) to be multi-homed for VPLS to multiple routers deployed as PEs without requiring the use of mVPLS.

In the configuration example displayed in Figure 41, the MTUs have spoke SDPs to two PEs devices. One is designated as the primary and one as the secondary spoke SDP. This is based on a precedence value associated with each spoke. If the primary and secondary spoke-SDPs have the same precedence value, the spoke-SDP with lower ID functions as the primary SDP.

The secondary spoke is in a blocking state (both on receive and transmit) as long as the primary spoke is available. When the primary spoke becomes unavailable (due to link failure, PEs failure, etc.), the MTU immediately switches traffic to the backup spoke and starts receiving/sending traffic to/from the standby spoke. Optional revertive operation (with configurable switch-back delay) is applicable only when one of the spokes is configured with precedence of primary. If not, this action does not take place. Forced manual switchover is also supported.

To speed up the convergence time during a switchover, MAC flush is configured. The MTUs generates a MAC flush message over the newly unblocked spoke when a spoke change occurs. As a result, the PEs receiving the MAC flush will flush all MACs associated with the impacted VPLS service instance and forward the MAC flush to the other PEs in the VPLS network if “propagate-mac-flush” is enabled.

Inter-Domain VPLS Resiliency Using Multi-Chassis Endpoints

Note: MC-EP is not supported in 7210 SAS devices. This section serves to provide an example on how 7210 SAS devices can be used as a MTUs device in an MC-EP solution. In this solution the 7750 SR routers provide the MC-EP functionality.

Inter-domain VPLS refers to a VPLS deployment where sites may be located in different domains. An example of inter-domain deployment can be where different Metro domains are interconnected over a Wide Area Network (Metro1-WAN-Metro2) or where sites are located in different autonomous systems (AS1-ASBRs-AS2).

Multi-chassis endpoint (MC-EP) provides an alternate solution that does not require RSTP at the gateway VPLS PEs while still using pseudowires to interconnect the VPLS instances located in the two domains.

MC-EP expands the single chassis endpoint based on active-standby pseudowires for VPLS shown in Figure 42. In the solution depicted by the Figure 42, 7210 devices are used as MTUs.

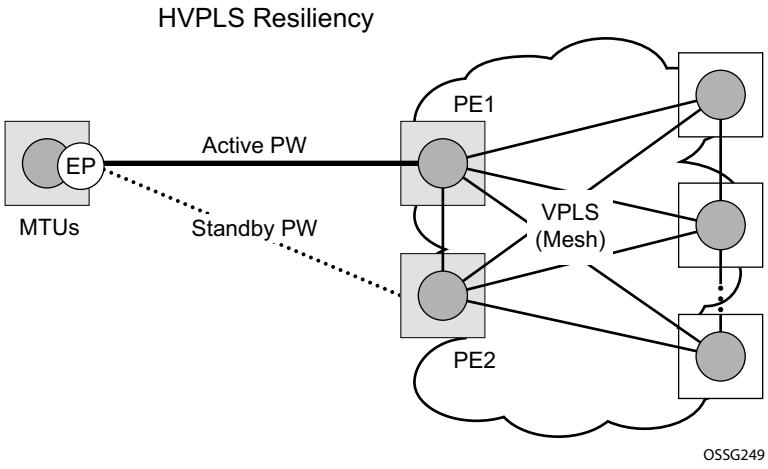

Figure 42: HVPLS Resiliency Based on AS Pseudowires

The active-standby pseudowire solution is appropriate for the scenario when only one VPLS PE (MTU-s) needs to be dual-homed to two core PEs (PE1 and PE2).

VPLS Access Redundancy

A second application of hierarchical VPLS is using MTUs that are MPLS-enabled which must have spoke SDPs to the closest PE node. To protect against failure of the PE node, an MTU can be dual-homed.

Listed below are several mechanisms that can be used to resolve a loop in an access network where 7210s are used

- STP-based access, with or without mVPLS.
- Ethernet APS using G.8032.

STP-Based Redundant Access to VPLS

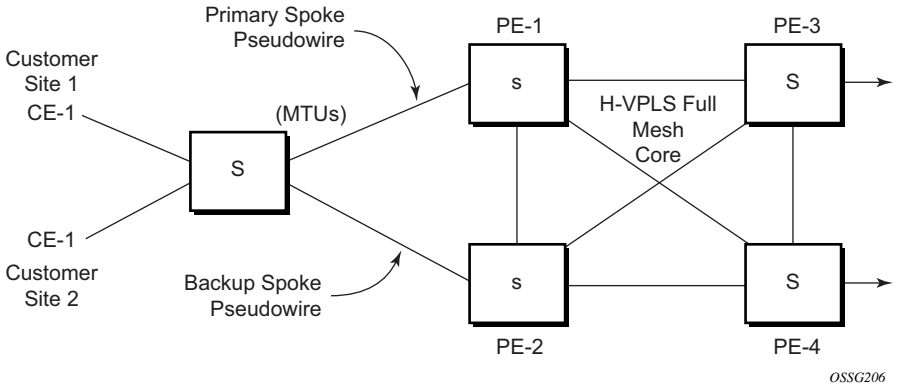

Figure 43: Dual Homed MTU-s in Two-Tier Hierarchy H-VPLS

In configuration shown in Figure 43, STP is activated on the MTU and two PEs in order to resolve a potential loop.

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. Running STP in every VPLS in this topology is not efficient as the node includes functionality which can associate a number of VPLSes to a single STP instance running over the redundant SDPs. Node redundancy is thus achieved by running STP in one VPLS. Thus, this applies the conclusions of this STP to the other VPLS services.

The VPLS instance running STP is referred to as the “management VPLS” or mVPLS. In the case of a failure of the active node, STP on the management VPLS in the standby node will change the link states from disabled to active. The standby node will then broadcast a MAC flush LDP control message in each of the protected VPLS instances, so that the address of the newly active node can

be re-learned by all PEs in the VPLS. It is possible to configure two management VPLS services, where both VPLS services have different active spokes (this is achieved by changing the path-cost in STP). By associating different user VPLSes with the two management VPLS services, load balancing across the spokes can be achieved.

In this configuration the scope of STP domain is limited to MTU and PEs, while any topology change needs to be propagated in the whole VPLS domain.

This is done by using “MAC-flush” messages defined by RFC 4762, *Virtual Private LAN Services Using LDP Signaling*. In the case where STP acts as a loop resolution mechanism, every Topology Change Notification (TCN) received in a context of STP instance is translated into an LDP-MAC address withdrawal message (also referred to as a MAC-flush message) requesting to clear all FDB entries except the ones learned from the originating PE. Such messages are sent to all PE peers connected through SDPs (mesh and spoke) in the context of VPLS service(s) which are managed by the given STP instance.

Redundant Access to VPLS Without STP

The Alcatel-Lucent implementation also alternative methods for providing a redundant access to LAYER 2 services, such as MC-LAG, MC-APS or MC-RING. Also in this case, the topology change event needs to be propagated into VPLS topology in order to provide fast convergence.

Figure 41 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 an alternative route exists.

Note that the message described here is different than the message described in previous section and in RFC 4762, *Virtual Private LAN Services Using LDP Signaling*. The difference is in the interpretation and action performed in the receiving PE. According to the standard definition, upon receipt of a MAC withdraw message, all MAC addresses, except the ones learned from the source PE, are flushed,

This section specifies that all MAC addresses learned from the source are flushed. This message has been implemented as an LDP address message with vendor-specific type, length, value (TLV), and is called the flush-mine message.

The advantage of this approach (as compared to RSTP based methods) is that only MAC-affected addresses are flushed and not the full forwarding database. While this method does not provide a mechanism to secure alternative loop-free topology, the convergence time is dependent on the speed of the given CE device will open alternative link (L2-B switch in Figure 57) as well as on the speed PE routers will flush their FDB.

In addition, this mechanism is effective only if PE and CE are directly connected (no hub or bridge) as it reacts to physical failure of the link.

MAC Flush Message Processing

The previous sections described operation principle of several redundancy mechanisms available in context of VPLS service. All of them rely on MAC flush message as a tool to propagate topology change in a context of the given VPLS. This section aims to summarize basic rules for generation and processing of these messages.

As described on respective sections, the 7210 SAS supports two types of MAC flush message, flush-all-but-mine and flush-mine. The main difference between these messages is the type of action they signal. Flush-all-but-mine requests clearing of all FDB entries which were learned from all other LDP peers except the originating PE. This type is also defined by RFC 4762 as an LDP MAC address withdrawal with an empty MAC address list.

Flush-all-mine message requests clearing all FDB entries learned from originating PE. This means that this message has exactly other effect then flush-all-but-mine message. This type is not included in RFC 4762 definition and it is implemented using vendor specific TLV.

The advantages and disadvantages of the individual types should be apparent from examples in the previous section. The description here focuses on summarizing actions taken on reception and conditions individual messages are generated.

Upon reception of MAC flush messages (regardless the type) SR-Series PE will take following actions:

- Clears FDB entries of all indicated VPLS services conforming the definition.
- Propagates the message (preserving the type) to all LDP peers, if “propagate-mac-flush” flag is enabled at corresponding VPLS level.

The flush-all-but-mine message is generated under following conditions:

- The flush-all-but-mine message is received from LDP peer and propagate-mac-flush flag is enabled. The message is sent to all LDP peers in the context of VPLS service it was received in.
- TCN message in a context of STP instance is received. The flush-all-but-mine message is sent to all LDP-peers connected with spoke and mesh SDPs in a context of VPLS service controlled by the given STP instance (based on mVPLS definition). The message is sent only to LDP peers which are not part of STP domain, which means corresponding spoke and mesh SDPs are not part of mVPLS.
- Flush-all-but-mine message is generated when switch over between spoke SDPs of the same endpoint occurs. The message is sent to LDP peer connected through newly active spoke SDP.

The flush-mine message is generated under following conditions:

- The flush-mine message is received from LDP peer and “propagate-mac-flush” flag is enabled. The message is sent to all LDP peers in the context of VPLS service it was received.
- The flush-mine message is generated when on a SAP or SDP transition from operationally up to an operationally down state and send-flush-on-failure flag is enabled in the context of the given VPLS service. The message is sent to all LDP peers connected in the context of the given VPLS service. Note, that enabling “send-flush-on-failure” the flag is blocked in VPLS service managed by mVPLS. This is to prevent that both messages are sent at the same time.
- The flush-mine message is generated when on a MC-LAG SAP or MC-APS SAP transition from an operationally up state to an operationally down state. The message is sent to all LDP peers connected in the context of the given VPLS service.
- The flush-mine message is generated when on a MC-RING SAP transition from operationally up to an operationally down state or when MC-RING SAP transitions to slave state. The message is sent to all LDP peers connected in the context of the given VPLS service.

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 44 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.

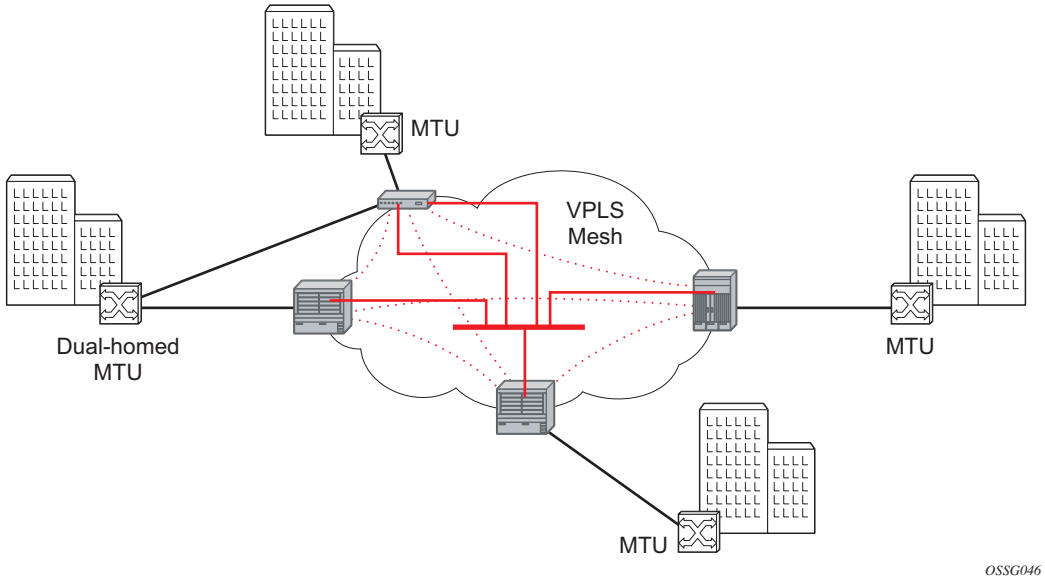

Figure 44: HVPLS with SAP Redundancy

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 7210 SAS M 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

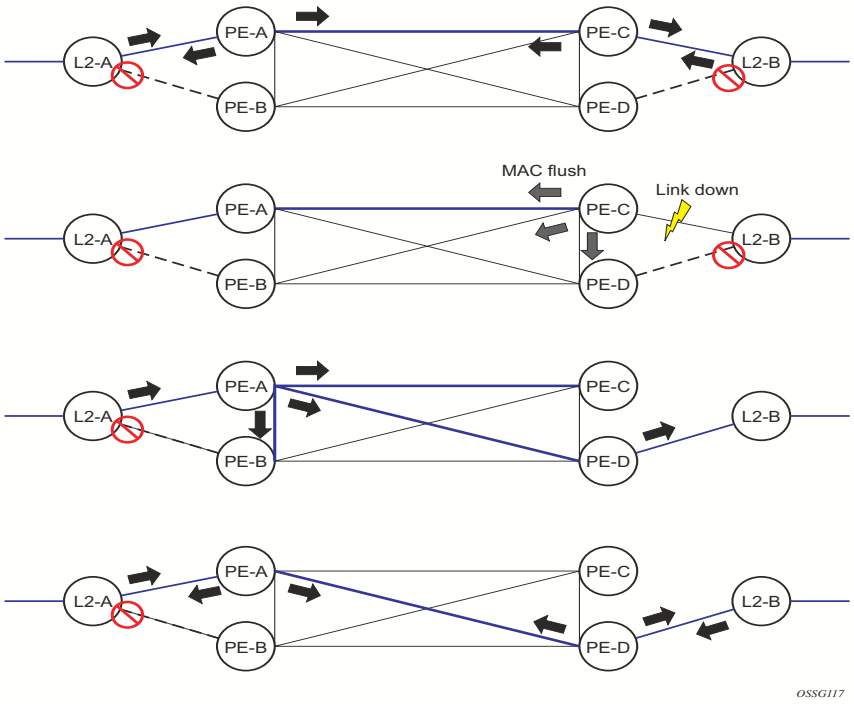

Figure 45: Dual Homed CE Connection to VPLS

Figure 45 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 an 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 45) 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.

VPLS Service Considerations

This section describes various 7210 SAS service features and any special capabilities or considerations as they relate to VPLS services.

SAP Encapsulations

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

- Ethernet null
- Ethernet Dot1q
- Ethernet Dot1q Default
- Ethernet Dot1q Explicit Null

VLAN Processing

The SAP encapsulation definition on Ethernet ingress ports defines which VLAN tags are used to determine the service that the packet belongs:

1. Null encapsulation defined on ingress — Any VLAN tags are ignored and the packet goes to a default service for the SAP.
2. Dot1q encapsulation defined on ingress — Only first label is considered.
3. Dot1q Default encapsulation defined on ingress — Tagged packets not matching any of the configured VLAN encapsulations would be accepted. This is like a default SAP for tagged packets.
4. Dot1q Explicit Null encapsulation defined on ingress — Any untagged or priority tagged packets will be accepted.

BGP Auto-Discovery for LDP VPLS

BGP Auto Discovery (BGP AD) for LDP VPLS is a framework for automatically discovering the endpoints of a Layer 2 VPN offering an operational model similar to that of an IP VPN. This model allows carriers to leverage existing network elements and functions, including but not limited to, route reflectors and BGP policies to control the VPLS topology.

BGP AD is an excellent complement to an already established and well deployed Layer 2 VPN signaling mechanism target LDP providing one touch provisioning for LDP VPLS where all the related PEs are discovered automatically. The service provider may make use of existing BGP policies to regulate the exchanges between PEs in the same, or in different, autonomous system (AS) domains. The addition of BGP AD procedures does not require carriers to uproot their existing VPLS deployments and to change the signaling protocol.

BGP AD Overview

The BGP protocol establishes neighbor relationships between configured peers. An open message is sent after the completion of the three-way TCP handshake. This open message contains information about the BGP peer sending the message. This message contains Autonomous System Number (ASN), BGP version, timer information and operational parameters, including capabilities. The capabilities of a peer are exchanged using two numerical values: the Address Family Identifier (AFI) and Subsequent Address Family Identifier (SAFI). These numbers are allocated by the Internet Assigned Numbers Authority (IANA). BGP AD uses AFI 65 (L2VPN) and SAFI 25 (BGP VPLS).

Information Model

Following is the establishment of the peer relationship, the discovery process begins as soon as a new VPLS service instance is provisioned on the PE.

Two VPLS identifiers are used to indicate the VPLS membership and the individual VPLS instance:

- VPLS-ID — Membership information, unique network wide identifier; same value assigned for all VPLS switch instances (VSIs) belonging to the same VPLS; encodable and carried as a BGP extended community in one of the following formats:
  - → A two-octet AS specific extended community
  - → An IPv4 address specific extended community

- VSI-ID— The unique identifier for each individual VSI, built by concatenating a route distinguisher (RD) with a 4 bytes identifier (usually the system IP of the VPLS PE); encoded and carried in the corresponding BGP NLRI.

In order to advertise this information, BGP AD employs a simplified version of the BGP VPLS NLRI where just the RD and the next 4 bytes are used to identify the VPLS instance. There is no need for Label Block and Label Size fields as T-LDP will take care of signaling the service labels later on.

The format of the BGP AD NLRI is very similar with the one used for IP VPN as depicted in Figure 46. The system IP may be used for the last 4 bytes of the VSI ID further simplifying the addressing and the provisioning process.

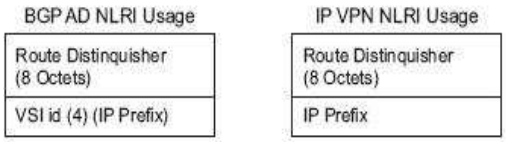

Figure 46: BGP AD NLRI versus IP VPN NLRI

Network Layer Reachability Information (NLRI) is exchanged between BGP peers indicating how to reach prefixes. The NLRI is used in the Layer 2 VPN case to tell PE peers how to reach the VSI rather than specific prefixes. The advertisement includes the BGP next hop and a route target (RT). The BGP next hop indicates the VSI location and is used in the next step to determine which signaling session is used for pseudowire signaling. The RT, also coded as an extended community, can be used to build a VPLS full mesh or a HVPLS hierarchy through the use of BGP import or export policies.

BGP is only used to discover VPN endpoints and the corresponding far end PEs. It is not used to signal the pseudowire labels. This task remains the responsibility of targeted-LDP (T-LDP).

FEC Element for T-LDP Signaling

Two LDP FEC elements are defined in RFC 4447, PW Setup & Maintenance Using LDP. The original pseudowire-ID FEC element 128 (0x80) employs a 32-bit field to identify the virtual circuit ID and it was used extensively in the initial VPWS and VPLS deployments. The simple format is easy to understand but it does not provide the required information model for BGP autodiscovery function. In order to support BGP AD and other new applications a new Layer 2 FEC element, the generalized FEC (0x81) is required.

The generalized pseudowire-ID FEC element has been designed for auto discovery applications. It provides a field, the address group identifier (AGI), that is used to signal the membership information from the VPLS-ID. Separate address fields are provided for the source and target address associated with the VPLS endpoints called the Source Attachment Individual Identifier (SAII) and respectively, Target Attachment Individual Identifier (TAII). These fields carry the VSI ID values for the two instances that are to be connected through the signaled pseudowire.

The detailed format for FEC 129 is depicted in Figure 47.

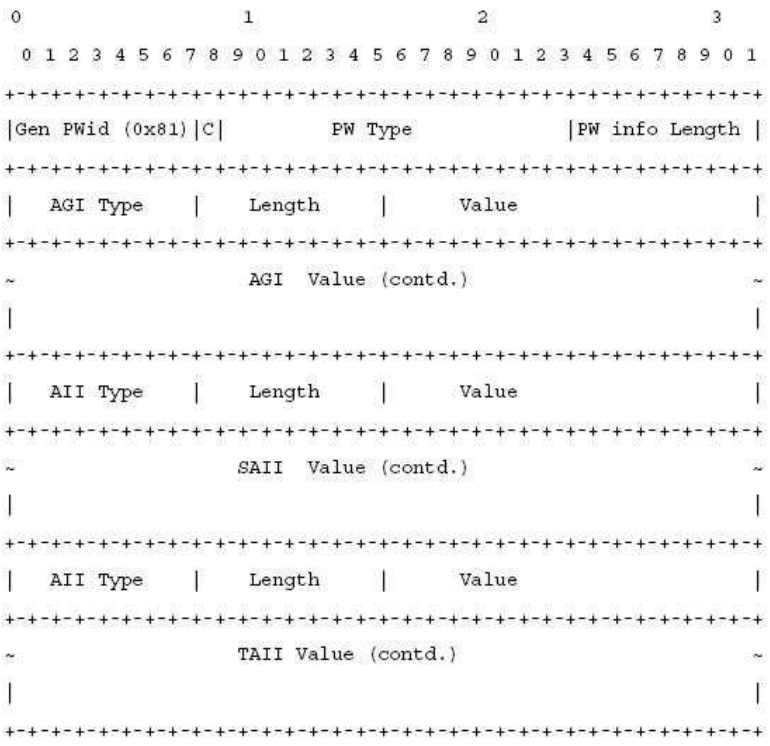

Figure 47: Generalized Pseudowire-ID FEC Element

Each of the FEC fields are designed as a sub-TLV equipped with its own type and length providing support for new applications. To accommodate the BGP AD information model the following FEC formats are used:

- AGI (type 1) is identical in format and content with the BGP extended community attribute used to carry the VPLS-ID value.
- Source AII (type 1) is a 4 bytes value destined to carry the local VSI-id (outgoing NLRI minus the RD).

- Target AII (type 1) is a 4 bytes value destined to carry the remote VSI-ID (incoming NLRI minus the RD).

BGP-AD and Target LDP (T-LDP) Interaction

BGP is responsible for discovering the location of VSIs that share the same VPLS membership. LDP protocol is responsible for setting up the pseudowire infrastructure between the related VSIs by exchanging service specific labels between them.

Once the local VPLS information is provisioned in the local PE, the related PEs participating in the same VPLS are identified through BGP AD exchanges. A list of far-end PEs is generated and triggers the creation, if required, of the necessary T-LDP sessions to these PEs and the exchange of the service specific VPN labels. The steps for the BGP AD discovery process and LDP session establishment and label exchange are shown in Figure 48.

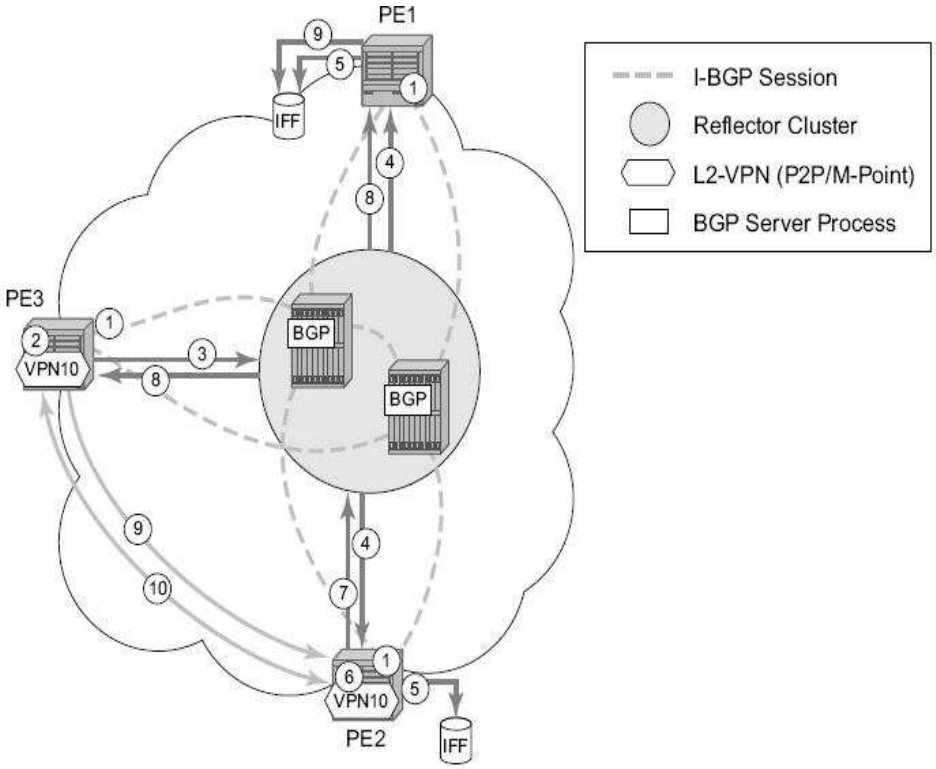

Figure 48: BGP-AD and T-LDP Interaction

Key:

1. Establish I-BGP connectivity RR.
2. Configure VPN (10) on edge node (PE3).

3. Announce VPN to RR using BGP-AD.
4. Send membership update to each client of the cluster.
5. LDP exchange or inbound FEC filtering (IFF) of non-match or VPLS down.
6. Configure VPN (10) on edge node (PE2).
7. Announce VPN to RR using BGP-AD.
8. Send membership update to each client of the cluster.
9. LDP exchange or inbound FEC filtering (IFF) of non-match or VPLS down.
10. Complete LDP bidirectional pseudowire establishment FEC 129.

SDP Usage

Service Access Points (SAP) are linked to transport tunnels using Service Distribution Points (SDP). The service architecture of the 7210 platform allows services to be abstracted from the transport network.

MPLS transport tunnels are signaled using the Resource Reservation Protocol (RSVP-TE) or by the Label Distribution Protocol (LDP). The capability to automatically create an SDP only exists for LDP based transport tunnels. Using a manually provisioned SDP is available for both RSVP-TE and LDP transport tunnels. Refer to the appropriate 7210 SAS OS MPLS Guide for more information about MPLS, LDP, and RSVP.

Automatic Creation of SDPs

When BGP AD is used for LDP VPLS and LDP is used as the transport tunnel there is no requirement to manually create an SDP. The LDP SDP can be automatically instantiated using the information advertised by BGP AD. This simplifies the configuration on the service node.

Enabling LDP on the IP interfaces connecting all nodes between the ingress and the egress, builds transport tunnels based on the best IGP path. LDP bindings are automatically built and stored in the hardware. These entries contain an MPLS label pointing to the best next hop along the best path toward the destination.

When two endpoints need to connect and no SDP exists, a new SDP will automatically be constructed. New services added between two endpoints that already have an automatically created SDP will be immediately used. No new SDP will be constructed. The far-end information is gleaned from the BGP next hop information in the NLRI. When services are withdrawn with a BGP_Unreach_NLRI, the automatically established SDP will remain up as long as at least one service is connected between those endpoints. An automatically created SDP will be removed and the resources released when the only or last service is removed.

Manually Provisioned SDP

The carrier is required to manually provision the SDP if they create transport tunnels using RSVP-TE. Operators have the option to choose a manually configured SDP, if they use LDP as the tunnel signaling protocol. The functionality is the same regardless of the signaling protocol.

Creating a BGP-AD enabled VPLS service on an ingress node with the manually provisioned SDP option causes the Tunnel Manager to search for an existing SDP that connects to the far-end PE. The far-end IP information is gleaned from the BGP next hop information in the NLRI. If a single SDP exists to that PE, it is used. If no SDP is established between the two endpoints, the service remains down until a manually configured SDP becomes active.

When multiple SDPs exist between two endpoints, the tunnel manager selects the appropriate SDP. The algorithm preferred SDPs with the best (lower) metric. Should there be multiple SDPs with equal metrics, the operational state of the SDPs with the best metric is considered. If the operational state is the same, the SDP with the higher sdp-id is used. If an SDP with a preferred metric is found with an operational state that is not active, the tunnel manager flags it as ineligible and restarts the algorithm.

Automatic Instantiation of Pseudowires (SDP Bindings)

The choice of manual or auto provisioned SDPs has limited impact on the amount of required provisioning. Most of the savings are achieved through the automatic instantiation of the pseudowire infrastructure (SDP bindings). This is achieved for every auto-discovered VSIs through the use of the pseudowire template concept. Each VPLS service that uses BGP AD contains the “pw-template-binding” option defining specific layer 2 VPN parameters. This command references a “pw-template” which defines the pseudowire parameters. The same “pwtemplate” may be referenced by multiple VPLS services. As a result, changes to these pseudowire templates have to be treated with great care as they may impact many customers at once.

The Alcatel-Lucent implementation provides for safe handling of pseudowire templates. Changes to the pseudowire templates are not automatically propagated. Tools are provided to evaluate and distribute the changes. The following command is used to distribute changes to a “pw-template” at the service level to one or all services that use that template.

PERs-4# tools perform service id 300 eval-pw-template 1 allow-service-impact

If the service ID is omitted, then all services are updated. The type of change made to the “pwtemplate” influences how the service is impacted.

1. Adding or removing a split-horizon-group will cause the router to destroy the original object and recreate using the new value.

2. Changing parameters in the vc-type {ether | vlan} command requires LDP to re-signal the labels.

Both of these changes affect the services. Other changes are not service affected.

Mixing Statically Configured and Auto-Discovered Pseudowires in a VPLS service

The services implementation allows for manually provisioned and auto-discovered pseudowire (SDP bindings) to co-exist in the same VPLS instance (for example, both FEC128 and FEC 129 are supported). This allows for gradual introduction of auto discovery into an existing VPLS deployment.

As FEC 128 and 129 represent different addressing schemes, it is important to make sure that only one is used at any point in time between the same two VPLS instances. Otherwise, both pseudowires may become active causing a loop that might adversely impact the correct functioning of the service. It is recommended that FEC128 pseudowire be disabled as soon as the FEC129 addressing scheme is introduced in a portion of the network. Alternatively, RSTP may be used during the migration as a safety mechanism to provide additional protection against operational errors.

Resiliency Schemes

The use of BGP-AD on the network side, or in the backbone, does not affect the different resiliency schemes Alcatel-Lucent has developed in the access network. This means that both Multi-Chassis Link Aggregation (MC-LAG) and Management-VPLS (M-VPLS) can still be used.

BGP-AD may co-exist with Hierarchical-VPLS (H-VPLS) resiliency schemes (for example, dual homed MTU-s devices to different PE-rs nodes) using existing methods (M-VPLS and statically configured Active or Standby pseudowire endpoint).

If provisioned SDPs are used by BGP AD, M-VPLS may be employed to provide loop avoidance. However, it is currently not possible to auto-discover active or standby pseudowires and to instantiate the related endpoint.

Routed VPLS

Routed VPLS (R-VPLS) allows a VPLS instance to be associated with an IES IP interface.

Note: RVPLS is supported on 7210 SAS-M network mode (and not in 7210 SAS-T network mode) In network mode R-VPLS service is available to provide customer service and inband management of the node. In access-uplink, it is supported only for inband management of the device on both 7210 SAS-M and 7210 SAS-T.

Within an R-VPLS service, traffic with a destination MAC matching that of the associated IP interface is routed based on the IP forwarding table; all other traffic is forwarded based on the VPLS forwarding table.

The R-VPLS service can be associated with an IPv4 interface and supports only static routing. It is primarily designed for use of inband management of the node when operating the node in Access-Uplink mode. It is useful for an inband management of ring for 7210 nodes using a single IPv4 subnet.

IES IP Interface Binding

A standard IP interface within an existing IES service context may be bound to a service name. A VPLS service only supports binding for a single IP interface.

While an IP interface may only be bound to a single VPLS service, the routing context containing the IP interface (IES) may have other IP interfaces bound to other VPLS service contexts. In other words, Routed VPLS allows the binding of IP interfaces in IES services to be bound to VPLS services.

Assigning a Service Name to a VPLS Service

When a service name is applied to any service context, the name and service ID association is registered with the system. A service name cannot be assigned to more than one service ID. Special consideration is given to a service name that is assigned to a VPLS service that has the "*configure>service>vpls>allow-ip-int-binding*" command is enabled. If a name is applied to the VPLS service while the flag is set, the system scans the existing IES services for an IP interface that is bound to the specified service name. If an IP interface is found, the IP interface is attached to the VPLS service associated with the name. Only one interface can be bound to the specified name.

If the allow-ip-int-binding command is not enabled on the VPLS service, the system does not attempt to resolve the VPLS service name to an IP interface. As soon as the allow-ip-int-binding

flag is configured on the VPLS, the corresponding IP interface is adhered and become operational up. There is no need to toggle the shutdown or no shutdown command.

If an IP interface is not currently bound to the service name used by the VPLS service, no action is taken at the time of the service name assignment.

Service Binding Requirements

In the event that the defined service name is created on the system, the system checks to ensure that the service type is VPLS. If the created service type is VPLS, the IP interface is eligible to enter the operationally upstate.

Bound Service Name Assignment

In the event that a bound service name is assigned to a service within the system, the system first checks to ensure the service type is VPLS. Secondly the system ensures that the service is not already bound to another IP interface through the service name. If the service type is not VPLS or theservice is already bound to another IP interface through the service ID, the service name assignment fails.

A single VPLS instance cannot be bound to two separate IP interfaces.

Binding a Service Name to an IP Interface

An IP interface within an IES service context may be bound to a service name at anytime. Only one interface can be bound to a service. When an IP interface is bound to a service name and the IP interface is administratively up, the system scans for a VPLS service context using the name and takes the following actions:

- If the name is not currently in use by a service, the IP interface is placed in an operationally down: Non-existent service name or inappropriate service type state.
- If the name is currently in use by a non-VPLS service or the wrong type of VPLS service, the IP interface is placed in the operationally down: Non-existent service name or inappropriate service type state.
- If the name is currently in use by a VPLS service without the allow-ip-int-binding flag set, the IP interface is placed in the operationally down: VPLS service allow-ip-intbinding flag not set state. There is no need to toggle the shutdown or no shutdown command.
- If the name is currently in use by a valid VPLS service and the allow-ip-int-binding flag is set, the IP interface is eligible to be placed in the operationally up state depending on other operational criteria being met.

IP Interface Attached VPLS Service Constraints

Once a VPLS service has been bound to an IP interface through its service name, the service name assigned to the service cannot be removed or changed unless the IP interface is first unbound from the VPLS service name.

A VPLS service that is currently attached to an IP interface cannot be deleted from the system unless the IP interface is unbound from the VPLS service name.

The allow-ip-int-binding flag within an IP interface attached VPLS service cannot be reset. The IP interface must first be unbound from the VPLS service name to reset the flag.

IP Interface and VPLS Operational State Coordination

When the IP interface is successfully attached to a VPLS service, the operational state of the IP interface is dependent upon the operational state of the VPLS service.

The VPLS service itself remains down until at least one virtual port (SAP, spoke-SDP or Mesh-SDP) is operational.

IP Interface MTU and Fragmentation

In 7210 SAS-M and 7210 SAS-T Access-Uplink mode, VPLS service MTU is not supported. The user must ensure that the port MTU is configured appropriately so that the largest packet traversing through any of the SAPs (virtual ports) of the VPLS service can be forwarded out of any of the SAPs. VPLS services do not support fragmentation and can discard packets larger than the configured port MTU.

When an IP interface is associated with a VPLS service, the IP-MTU is based on either the administrative value configured for the IP interface or an operational value derived from port MTU of all the SAPs configured in the service. The port MTU excluding the Layer 2 Header and tags for all the ports which have SAPs configured in this VPLS service are considered and the minimum value among those are computed (which is called computed MTU). The operational value of the IP interface is set as follows:

- If the configured (administrative) value of IP MTU is greater than the computed MTU, then the operational IP MTU is set to the computed MTU.
- If the configured (administrative) value of IP MTU is lesser than or equal to the computed MTU, then operational IP MTU is set to the configured (administrative) value of IP MTU.

Unicast IP Routing into a VPLS Service

The IP interface MTU parameters may be changed at anytime.

ARP and VPLS FIB Interactions

Two address-oriented table entries are used when routing into a VPLS service. On the routing side, an ARP entry is used to determine the destination MAC address used by an IP next-hop. In the case where the destination IP address in the routed packet is a host on the local subnet represented by the VPLS instance, the destination IP address itself is used as the next-hop IP address in the ARP cache lookup. If the destination IP address is in a remote subnet that is reached by another router attached to the VPLS service, the routing lookup returns the local IP address on the VPLS service of the remote router is returned. If the next-hop is not currently in the ARP cache, the system generates an ARP request to determine the destination MAC address associated with the next-hop IP address. IP routing to all destination hosts associated with the next-hop IP address stops until the ARP cache is populated with an entry for the next-hop. The dynamically populated ARP entries age out according to the ARP aging timer.

NOTE: In 7210 static ARP, entries cannot be used.

The second address table entry that affects VPLS routed packets is the MAC destination lookup in the VPLS service context. The MAC associated with the ARP table entry for the IP next-hop may or may not currently be populated in the VPLS Layer 2FIB table. While the destination MAC is unknown (not populated in the VPLS FIB), the system is flooded with all packets destined to that MAC (routed or bridged) to all virtual ports within the VPLS service context. Once the MAC is known (populated in the VPLS FIB), all packets destined to the MAC (routed or bridged) is targeted to the specific virtual port where the MAC has been learned. As with ARP entries, static MAC entries may be created in the VPLS FIB. Dynamically learned MAC addresses are allowed to age out or be flushed from the VPLS FIB while static MAC entries always remain associated with a specific virtual port. Dynamic MACs may also be relearned on another VPLS virtual port than the current virtual port in the FIB. In this case, the system automatically moves the MAC FIB entry to the new VPLS virtual port.

NOTE: In 7210 SAS, whenever a MAC entry is removed from the VPLS FIB (either explicitly by the user or due to MAC aging or mac-move), ARP entries which match this MAC address is removed from the ARP cache. Though the VPLS FIB entries are not removed; an ARP entry ages out and is removed from the ARP cache.

NOTE: If the VPLS FIB limit is reached and we are no longer able to learn new MAC address, ARP will also not be learnt. In network mode, RVPLS SAPs cannot be configured on Hybrid port.

Routed VPLS Specific ARP Cache Behavior

In typical routing behavior, the system uses the IP route table to select the egress interface and then at the egress forwarding engine, an ARP entry is used forward the packet to the appropriate Ethernet MAC. With routed VPLS, the egress IP interface may be represented by multiple egress (VPLS service virtual ports).

The following tables describes how the ARP cache and MAC FIB entry states interact.

Table 18: Routing behavior in RVPLS and interaction ARP Cache and MAC FIB

ARP Cache Entry	MAC FIB Entry	Routing or System behavior
ARP Cache Miss (No Entry)	Known or Unknown	Triggers a request to control plane ARP processing module, to send out an ARP request, out of all the SAPs. (also known as virtual ports) of the VPLS instance.
ARP Cache Hit	Known	Forward to specific VPLS virtual port or SAP.
	Unknown	This behavior cannot happen typically in 7210 SAS, as and when a L2 entry is removed from the FDB, the matching MAC address is also removed from the ARP cache.

The allow-ip-int-binding VPLS Flag

The allow-ip-int-binding flag on a VPLS service context is used to inform the system that the VPLS service is enabled for routing support. The system uses the setting of the flag as a key to determine what type of ports the VPLS service may span.

The system also uses the flag state to define which VPLS features are configurable on the VPLS service to prevent enabling a feature that is not supported when routing support is enabled.

Routed VPLS SAPs only Supported on Standard Ethernet Ports

The allow-ip-int-binding flag is set (routing support enabled) on a VPLS service. SAPs within the service can be created on standard Ethernet ports.

LAG Port Membership Constraints

If a LAG has a non-supported port type as a member, a SAP for the routing-enabled VPLS service cannot be created on the LAG. Once one or more routing enabled VPLS SAPs are associated with a LAG, a non-supported Ethernet port type cannot be added to the LAG membership.

VPLS Feature Support and Restrictions

When the allow-ip-int-binding flag is set on a VPLS service, the following features cannot be enabled (The flag also cannot be enabled while any of these features are applied to the VPLS service). The following restrictions apply to both network mode and access-uplink moed unless called out separately:

- SDPs used in spoke or mesh SDP bindings cannot be configured.
- In access-uplink mode, the VPLS service type cannot be M-VPLS.
- In network mode, the VPLS Service type must be 'r-vpls' and any other VPLS service is not allowed.
- MVR from Routed VPLS and to another SAP is not supported.
- Default QinQ SAPs is not supported in R-VPLS service.
- The “allow-ip-int-binding” command cannot be used in a VPLS service which is acting as the G8032 control instance.
- IPv4 filters (ingress and egress) can be used with the R-VPLS SAPs. Additionally IP ingress override filters are supported which affects the behavior of the IP filters attached to the R-VPLS SAPs. Please see below for more information about use of ingress override filters.
- MAC filters (ingress and egress) are not supported for use with R-VPLS SAPs.
- VPLS IP interface is not allowed in a R-VPLS service. The converse also holds.
- In network mode, VPLS service can use the following 'svc-sap-type' values: any.
- In Access-uplink mode, the VPLS service can use either access SAP or Access-Uplink SAPs. In network mode, the VPLS service can use only access SAPs.

- In Access-uplink mode, VPLS service can use the following 'svc-sap-type' values: any, dot1q-preserve and null-star. Only specific SAP combinations are allowed for a given svc-sap-type, except that default QinQ SAPs cannot be used in a R-VPLS service. The allowed SAP combinations are similar to that available in a plain VPLS service and is as given in the table above in the services Chapter (with the exception noted before).
- G8032 or mVPLS/STP based protection mechanism can be used with R-VPLS service. A separate G8032 control instance or a separate mVPLS/STP instance needs to be used and the R-VPLS SAPs needs to be associated with these control instances such that the R-VPLS SAP's forwarding state is driven by the control instance protocols.
- IGMP snooping is not supported in a VPLS service. IP multicast is not supported in the R-VPLS service.
- DHCP snooping is not supported for the SAPs configured in the routed VPLS service. Instead, DHCP relay can be enabled on the IES service associated with the routed VPLS service.

VPLS SAP Ingress IP Filter Override

When an IP Interface is attached to a VPLS service context, the VPLS SAP provisioned IP filter for ingress routed packets may be optionally overridden in order to provide special ingress filtering for routed packets. This allows different filtering for routed packets and non-routed packets. The filter override is defined on the IP interface bound to the VPLS service name. A separate override filter may be specified for IPv4 packet types.

If a filter for a given packet type (IPv4) is not overridden, the SAP specified filter is applied to the packet (if defined).

The following tables lists ACL Lookup behavior with and without Ingress Override filter attached to an IES interface in a R-VPLS service:

Table 19: ACL Lookup behavior with Ingress Override filter attached to an IES interface in a R-VPLS service.

Type of traffic	SAP Ingress IPv4 Filter	SAP Egress IPv4 Filter	Ingress Override IPv4 Filter
Destination MAC != IES IP interface MAC	Yes	Yes	No
Destination MAC = IES IP interface MAC and Destination IP on same subnet as IES interface	No	No	Yes

Table 19: ACL Lookup behavior with Ingress Override filter attached to an IES interface in a R-VPLS service.

Type of traffic	SAP Ingress IPv4 Filter	SAP Egress IPv4 Filter	Ingress Override IPv4 Filter
Destination Mac = IES IP interface mac and destination IP not on same subnet as IES IP interface and route to destination IP does not exist	No	No	No
Destination Mac = IES IP interface mac and destination IP not on same subnet as IES IP interface and route to destination IP exists	No	No	Yes
Destination MAC = IES IP interface MAC and IP TTL = 1	No	No	No
Destination MAC = IES IP interface MAC and IPv4 packet with Options	No	No	No
Destination MAC = IES IP interface MAC and IPv4 Multicast packet	No	No	No

Table 20: ACL Lookup behavior without Ingress Override filter attached to an IES interface in a R-VPLS service

Type of traffic	SAP Ingress IPv4 Filter	SAP Egress IPv4 Filter
Destination MAC != IES IP interface MAC	Yes	Yes
Destination MAC = IES IP interface MAC and Destination IP on same subnet as IES IP interface	Yes	No

Table 20: ACL Lookup behavior without Ingress Override filter attached to an IES interface in a R-VPLS service

Type of traffic	SAP Ingress IPv4 Filter	SAP Egress IPv4 Filter
Destination Mac = IES IP interface mac and destination IP not on same subnet as IES IP interface and route to destination IP does not exist	No	No
Destination Mac = IES IP interface MAC and destination IP not on same subnet as IES IP interface and route to destination IP exists	Yes	No
Destination MAC = IES IP interface MAC and IP TTL = 1	No	No
Destination MAC = IES IP interface MAC and IPv4 packet with Options	No	No
Destination MAC = IES IP interface MAC and IPv4 Multicast packet	No	No

QoS Support for VPLS SAPs and IP interface in a Routed VPLS service

- SAP ingress classification (IPv4 and MAC criteria) is supported for SAPs configured in the service. SAP ingress policies cannot be associated with IES IP interface.
- Egress Port based queuing and shaping are available. It is shared among all the SAPs on the port.
- Port based Egress Marking is supported for both routed packets and bridged packets. The existing access egress QoS policy can be used for Dot1p marking.
- In Access-Uplink mode, IES IP interface bound to routed VPLS services, IES IP interface on access SAPs and IES IP interface on Access-Uplink SAPs are designed for use with inband management of the node. Consequently, they share a common set of queues for CPU bound management traffic. All CPU bound traffic is policed to pre-defined rates before being queued into CPU queues for application processing. The system uses meters per application or a set of applications. It does not allocate meters per IP interface. The possibility of CPU overloading has been reduced by use of these mechanisms. Users must use appropriate security policies either on the node or in the network to ensure that this does not happen.

Routed VPLS Supported Routing Related Protocols

In network mode and access-uplink mode R-VPLS is supported only in the base routing instance. Only IPv4 addressing support is available for IES interfaces associated with Routed VPLS service. The following lists the support available for routing protocols on IP interfaces bound to a VPLS service in access-uplink mode and network mode.

Table 21:

Services	Access-uplink	Network
Static-routing	Supported	Supported
BGP	Not Supported	Not Supported
OSPF	Not Supported	Supported
ISIS	Not Supported	Supported
BFD	Not Supported	Supported
VRRP	Not Supported	Supported
ARP and Proxy-Arp	ARP is supported	Both are supported
DHCP Relay (Note-1)	Supported	Supported

NOTE 1: DHCP relay can be configured for the IES interface associated with the Routed VPLS service. DHCP snooping cannot be configured on the VPLS SAPs in the routed VPLS Service.

Spanning Tree and Split Horizon

A routed VPLS context supports all spanning tree and port-based split horizon capabilities that a non-routed VPLS service supports.

Routed VPLS support available and Caveats

Routed VPLS supported functionality and restrictions for both access-uplink and network mode is given below. The following is applicable to both the modes, unless called out explicitly.

- Static ARP cannot be configured with an IES IP interface that is associated with an R-VPLS, though static MAC can be configured in an R-VPLS service.
- Only Static routes are supported. No dynamic routing protocols are supported.
- Whenever a VPLS FIB entry is removed either due to user action, aging or mac-move, the corresponding ARP entry whose MAC address matches that of the MAC in the FIB is removed from the ARP cache.
- In network mode and access-uplink mode R-VPLS is supported only in the base routing instance. Only IPv4 addressing support is available for IES interfaces associated with Routed VPLS service.
- IPv6 addressing support is not available for IES interface associated with R-VPLS service.
- In network mode, R-VPLS service cannot be bound to an VPRN Service.
- In both network mode and access-uplink mode, multiple SAPs configured on the same port cannot be part of the same R-VPLS Service. In other words, a single service can only be configured with a single SAP on a given port.
- Service MTU configuration is not supported in the R-VPLS service.
- In network mode, in 'any' service (that is, svc-sap-type set to any), null sap accepts only un-tagged packets. Tagged packets received are dropped.
- MPLS protocols (For example: RSVP, LDP) cannot be enabled on R-VPLS IP interface
- MPLS-TP cannot use R-VPLS, IES, and IP interface.
- Discard-unknown feature is not supported in the VPLS service associated with R-VPLS (only on 7210 SAS-X,R6)
- In network mode, R-VPLS SAPS and IES interface can be configured on a MC-LAG LAG.

Epipe Emulation using Dot1q VLAN range SAP in VPLS with G8032

NOTE: This feature is supported only on 7210 SAS-M (10G variant, ETR and non-ETR) in Access-uplink mode only and 7210 SAS-T (all variants) in Access-uplink mode (also known as L2 mode).

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

The Figure 49 shows how two business offices, served by an operator are connected in a ring network deployment using Dot1q range SAPs and a VPLS service with G.8032 for protection.

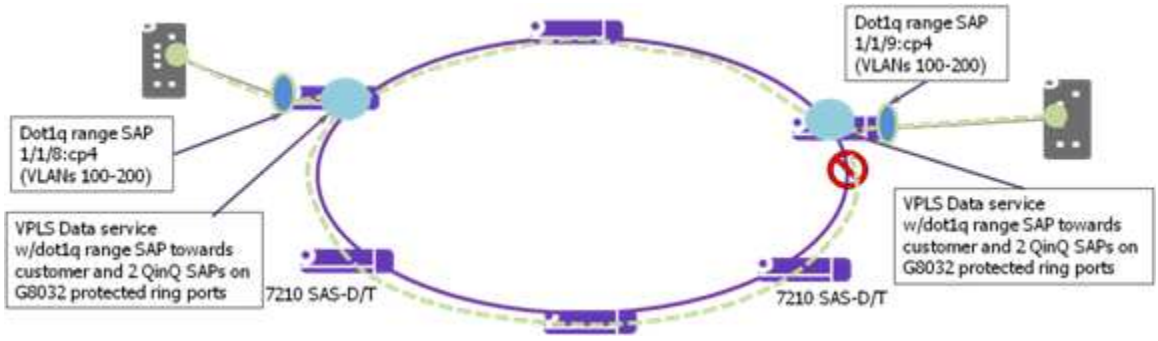

Figure 49: Epipe Emulation in a ring using VPLS with G.8032

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

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

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

The 7210 SAS platforms, currently support the use of Dot1q range SAPs with only Epipe services in either network/MPLS mode or access-uplink/L2 mode. Dot1q range SAPs allows operators to transport a range of VLANs by providing similar service treatment (service treatment refers to forwarding decision along with encapsulation used, QoS and ACL processing, accounting, etc.) to all the VLANs configured in the range. It simplifies service configuration and allows operators to scale the number of VLANs that can be handled by the node. This took care of the need to support hundreds of VLANs using a single SAP or a small number of SAPs. When MPLS the mode is deployed in ring topology, operators have the option of using different redundancy mechanisms such as FRR, primary/secondary LSPs, Active/Standby PWs, to improve Epipe service availability. No such option is available to protect Epipe service in L2 mode when deployed in a ring topology. Additionally many operators prefer G.8032 based ring protection mechanism, since a single control instance on the ring can potentially protect all the VPLS services on the ring.

This feature allows operators to deploy Epipe services in a ring topology when using L2 mode, by emulating an Epipe service using a VPLS service with G8032 protection and at the same time provides the benefits of using dot1q range SAPs. The user should ensure that the VPLS service is a point-to-point service. This is achieved by configuring a VPLS service with an access dot1q range SAP used at the customer handoff on one node in the ring and an access dot1q range SAP in a customer handoff of a VPLS service on another node (that is, at the other end of the Epipe), such that there are only two endpoints for the service in the network.

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

Configuration guidelines and restrictions

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

- The user can define access dot1q range SAPs, which specifies a group of VLANs which receive similar service treatment, that is, forwarding behavior, SAP ingress QoS treatment and SAP (behavior similar to that available in Epipe service) and allows it to be configured in a VPLS service.
  - → On the node, where the service originates, in addition to the access dot1q range SAP, the service should be configured with Q1.* SAPs on the two G.8032 ring ports. The access or access-uplink Q1.*SAPs can be used, but the access-uplink SAPs are recommended for use.The user cannot configure any other SAPs in the same VPLS service.

  - → There is no special configuration required on intermediate nodes, that is, the ring nodes which do not terminate or originate the service. The nodes should be configured for providing transit VPLS service and the VPLS service must use the same G8032 instance for protection as is used by the service on originating and terminating node.
  - → The Epipe service on 7210, currently does not check if the inner tag received on a Q1.* SAP is within the range of the configured VLANs. VPLS service too has the same behavior.
- Support for SAP Ingress QoS, Ingress and Egress ACLs, accounting, and other services, for dot1q range SAP configured in a VPLS service matches the support available in Epipe service.
- G.8032 mechanism is used for loop detection in ring network and service protection. A separate VPLS service representing the G.8032 control instance must be configured and the state should be associated with this service.
  - → Use of dot1q range SAPs to provide service on the interconnection node, in a G.8032 major-ring/sub-ring deployment, when using the virtual channel, is not supported. This restriction is not applicable when the interconnection node in a G8032 major-ring/sub-ring is configured without a virtual channel.
- mVPLS/xSTP support is available for use with Q1.* SAP on the ring ports to break the loop. This is a add-on to the G.8032 support.
- Broadcast, Unknown Unicast and Multicast (BUM) traffic is flooded in the service.
- Learning is enabled on the service by default, to avoid the need to flood the service traffic out of one of the ring ports, after network MAC addresses are learnt. The user has an option to disable learning per service. Learning enable/disable per SAP is not supported.
- MAC limiting is available per service. MAC limiting per SAP is not supported.
- CFM OAM is supported. The support for UP MEPs on the dot1q range SAP in the service to be used for fault management and performance management using the CFM/Y.1731 OAM tools is available.
  - → Only UP MEP is allowed to be configured only on the dot1q VLAN range SAPs. CFM/Y.1731 tools can be used for trouble shooting and performance measurements. User must pick a VLAN value from the range of VLANs configured for the dot1-range SAP using the CLI command config>eth-cfm>domain>association>bridge-identifier VLAN and enable the use of using the CLI command primary-vlan-enable under the MEP CLI context. It is used as the VLAN tag in the packet header for all the CFM/Y.1731 messages sent out in the context of the UP MEP.
  - → Down MEPs and MIPs are not allowed to be configured.
  - → Fault propagation is not supported with UP MEPs for dot1q range SAP in access-uplink mode.
- CFM support is not available for SAPs on the ring ports.
- IGMP snooping and MVR is not supported.

Configuring a VPLS Service with CLI

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

Topics in this section include:

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 69)
- For a local service, configure two SAPs, specifying local access ports and encapsulation values.
- For a distributed service, configure a SAP and an SDP (only for 7210 SAS devices in network mode) for each far-end node.

The following example displays a sample configuration of a local VPLS service on ALA-1.

For 7210 SAS devices configured in access-uplink mode:

```
*A:SAS-M-A0-2>config>service>vpls# info
----------------------------------------------
            stp
                shutdown
            exit
            sap 1/1/1:10.* create
                ingress
                    filter mac 1
                exit
            exit
            sap 1/1/2:10.* create
            exit
            no shutdown
----------------------------------------------
*A:SAS-M-A0-2>config>service>vpls#
*A:ALA-1>config>service>vpls# info
----------------------------------------------
...
        vpls 9001 customer 6 create
            description "Local VPLS"
            stp
                shutdown
            exit
            sap 1/2/2:0 create
                description "SAP for local service"
            exit
            sap 1/1/5:0 create
                description "SAP for local service"
            exit
            no shutdown
----------------------------------------------
*A:ALA-1>config>service>vpls#
*A:ALA-1>config>service# info
----------------------------------------------
...
       vpls 7 customer 7 create
            stp
                shutdown
            exit
            sap 1/1/21 create
```

```
            exit
            sap lag-1:700 create
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

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."
            stp
                shutdown
            exit
            sap 1/1/5:16 create
                description "VPLS SAP"
            exit
            spoke-sdp 2:22 create
            exit
        exit
...
----------------------------------------------
*A:ALA-1>config>service#

*A:ALA-2>config>service# info
----------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
            stp
                shutdown
            exit
            sap 1/1/5:16 create
                description "VPLS SAP"
            exit
            spoke-sdp 2:22 create
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-2>config>service#

*A:ALA-3>config>service# info
----------------------------------------------
...
        vpls 9000 customer 6 create
            description "This is a distributed VPLS."
            stp
                shutdown
            exit
            sap 1/1/3:33 create
```

```
                description "VPLS SAP"
            exit
            spoke-sdp 2:22 create
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-3>config>service#
```

Common Configuration Tasks

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

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)
3. Modify STP default parameters (optional) (see VPLS and Spanning Tree Protocol on page 283)
4. Enable service

Configuring VPLS Components

Use the CLI syntax displayed below to configure the following entities:

Creating a VPLS Service

Use the following CLI syntax to create a VPLS service:

CLI Syntax: config>service# vpls *service-id* [customer *customer-id*] [create][vpn vpn-id] [m-vpls] (for 7210 SAS-M and 7210 SAS-T in Network mode)

```
config>service# vpls service-id [customer customer-id] [create][ vpn vpn-id] [m-vpls] <service-id> [customer <customer-id>] [create] [vpn <vpn-id>] [m-vpls] [svc-sap-type {null-star|dot1q-preserve|any}] [customer-vid <vlan-id>] (for 7210 SAS-M in Access uplink mode)
description description-string
no shutdown
```

The following example displays a VPLS configuration:

```
*A:ALA-1>config>service>vpls# info
----------------------------------------------
...
        vpls 1000 customer 1 create
            description "This is a VPLS with NULL SAP"
            stp
                shutdown
            exit
            no shutdown
        exit
        vpls 2000 customer 6 create
            description "This is a Distributed VPLS with DOT1Q SAP"
            stp
                shutdown
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-1>config>service>vpls#
```

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 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

The following example displays mac-move information.

```
*A:ALA-1# show service id 6 all
....
*A:ALA-1#
-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id        : 1150               Mac Move          : Disabled
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Table Size        : 1000               Total Count       : 1000
Learned Count     : 1000               Static Count      : 0
Remote Age        : 900                Local Age         : 300
High WaterMark    : 95%                Low Watermark     : 90%
Mac Learning      : Enabl              Discard Unknown   : Dsabl
Mac Aging         : Enabl              Relearn Only      : True
===============================================================================
....
*A:ALA-1#
```

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 337
- Mode on page 338
- Bridge Priority on page 338
- Max Age on page 339
- Forward Delay on page 339
- Hello Time on page 340
- MST Instances on page 341
- MST Max Hops on page 341
- MST Name on page 341
- MST Revision on page 341

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 7210 SAS M. When STP on the VPLS is administratively enabled, but the administrative state of a SAP is down, BPDUs received on such a SAP are discarded.

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 7210 SAS M supports several variants of the Spanning Tree protocol:

- `rstp` — Rapid Spanning Tree Protocol (RSTP) compliant with IEEE 802.1D-2004 - default mode.
- `dot1w` — Compliant with IEEE 802.1w.
- `comp-dot1w` — Operation as in RSTP but backwards compatible with IEEE 802.1w (this mode was introduced for interoperability with some MTU types).
- `mstp` — Compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q REV/D5.0-09/2005. This mode of operation is only supported in an mVPLS.
- `pmstp` — Compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q REV/D3.0-04/2005 but with some changes to make it backwards compatible to 802.1Q 2003 edition and IEEE 802.1w.

See section Spanning Tree Operating Modes on page 283 for details on these modes.

CLI Syntax: `config>service>vpls` *`service-id`*`# stp`
`mode {rstp | comp-dot1w | dot1w | mstp|pmstp}`
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.

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 by the BPDUs.
The default value of **max-age** is 20. This parameter can be modified within a range of 6 to 40, limited by the standard STP parameter interaction 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 by 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 348).

For timer-based transitions, the 802.1D-2004 standard defines an internal variable forward-delay, which is used in calculating the default number of seconds that a SAP spends in the discarding and learning states when transitioning to the forwarding state. The value of the forward-delay variable depends on the STP operating mode of the VPLS instance:

- in `rstp` mode, but only when the SAP has not fallen back to legacy STP operation, the value configured by the **hello-time** command is used;
- in all other situations, the value configured by the **forward-delay** command is used.

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 339.

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 within 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 a VPLS SAP

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

Use the following CLI syntax to create:

- Local VPLS SAPs on page 342
- Distributed VPLS SAPs on page 343

Local VPLS SAPs

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

```
*A:ALA-1>config>service# info
----------------------------------------------
       vpls 1150 customer 1 create
            fdb-table-size 1000
            fdb-table-low-wmark 5
            fdb-table-high-wmark 80
            local-age 60
            stp
                shutdown
            exit
            sap 1/1/1:1155 create
            exit
            sap 1/1/2:1150 create
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-1>config>service#
```

Distributed VPLS SAPs

Note: Distributed VPLS service is not supported on 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

To configure a distributed VPLS service, you must configure service entities on originating and far-end nodes. You must use the same service ID on all ends (for example, create a VPLS service ID 9000 on ALA-1, ALA-2, and ALA-3). A distributed VPLS consists of a SAP on each participating node and an SDP bound to each participating node.

For SDP configuration information, see Configuring an SDP on page 71. For SDP binding information, see Configuring SDP Bindings on page 353.

The following example displays a configuration of VPLS SAPs configured for ALA-1, ALA-2, and ALA-3.

```
*A:ALA-3>config>service# info
----------------------------------------------
      vpls 1150 customer 1 create
            fdb-table-size 1000
            fdb-table-low-wmark 5
            fdb-table-high-wmark 80
            local-age 60
            stp
                shutdown
            exit
            sap 1/1/1:1155 create
            exit
            sap 1/1/2:1150 create
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-3>config>service#
```

Configuring Default QinQ SAPs to Pass all Traffic from Access to Access-uplink Port without any Tag Modifications

Note: Default QinQ SAPs are supported only on 7210 SAS-M and 7210 SAS-T devices configured in access-uplink mode.

The following example displays the VPLS SAP configuration of Default QinQ SAPs:

```
ALA-1>config>service# vpls 9 customer 1 svc-sap-type null-star create
            shutdown
            stp
                shutdown
            exit
            sap 1/1/5:*.* create
                statistics
                    ingress
                    exit
                exit
            exit
            sap 1/1/6:*.* create
                statistics
                    ingress
                    exit
                exit
            exit
        exit
```

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.

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 346 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.

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 7210 SAS M 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 339). 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 347).

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 by a shared media, their SAPs should all be configured as shared, and timer-based transitions are used.

Valid values for SAP link-type are shared and pt-pt with pt-pt being the default.

CLI Syntax: `config>service>vpls>sap>stp#`
`link-type {pt-pt|shared}`
Default: `link-type pt-pt`
Restore Default: `no link-type`

MST Instances

The SAP mst-instance command is used to create MST instances at the SAP level. MST instance at a SAP level can be created only if MST instances are defined at the service level.

The parameters that can be defined per instance are mst-path-cost and mst-port-priority.

- mst-path-cost — Specifies path-cost within a given MST instance. The path-cost is proportional to link speed.
- mst-port-priority — Specifies the port priority within a given MST instance.

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 349
- Operationally Discarding on page 349
- Operationally Learning on page 349
- Operationally Forwarding on page 350

Operationally Disabled

Operationally disabled is the normal operational state for STP on a SAP in a VPLS that has any of the following conditions:

- VPLS state administratively down
- SAP state administratively down
- SAP state operationally down

If the SAP enters the operationally up state with the STP administratively up and the SAP STP state is up, the SAP will transition to the STP SAP discarding state.

When, during normal operation, the router detects a downstream loop behind a SAP, BPDUs can be received at a very high rate. To recover from this situation, STP will transition the SAP to disabled state for the configured forward-delay duration.

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 339.

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. The STP is associated with a VPLS service like PVST is per VLAN. The difference between the two encapsulations is in the Ethernet and LLC framing and a type-length-value (TLV) field trailing the BPDU.The encapsulation format cannot be configured by the user,the system automatically determines the encapsulation format based on the BPDUs received on the port.

The following table shows differences between Dot1d and PVST Ethernet BPDU encapsulations based on the interface encap-type field:

Table 22: SAP BPDU Encapsulation States

Field	dot1d encap-type null	dot1d encap-type dot1q	PVST encap-type null	PVST encap-type dot1q
Destination MAC	01:80:c2:00:00:00	01:80:c2:00:00:00	N/A	01:00:0c:cc:cc:cd
Source MAC	Sending Port MAC	Sending Port MAC	N/A	Sending Port MAC
EtherType	N/A	0x81 00	N/A	0x81 00
Dot1p and CFI	N/A	0xe	N/A	0xe
Dot1q	N/A	VPLS SAP ID	N/A	VPLS SAP encap value
Length	LLC Length	LLC Length	N/A	LLC Length
LLC DSAP SSAP	0x4242	0x4242	N/A	0xaaaa (SNAP)
LLC CNTL	0x03	0x03	N/A	0x03
SNAP OUI	N/A	N/A	N/A	00 00 0c (Cisco OUI)
SNAP PID	N/A	N/A	N/A	01 0b
CONFIG	Standard 802.1d	Standard 802.1d	N/A	Standard 802.1d
TLV: Type & Len	N/A	N/A	N/A	58 00 00 00 02

Table 22: SAP BPDU Encapsulation States (Continued)

TLV: VLAN	N/A	N/A	N/A	VPLS SAP encap value
Padding	As Required	As Required	N/A	As Required

Each SAP has a Read-Only operational state that shows which BPDU encapsulation is currently active on the SAP. The states are:

- Dot1d — This state specifies that the switch is currently sending IEEE 802.1d standard BPDUs. The BPDUs are tagged or non-tagged based on the encapsulation type of the egress interface and the encapsulation value defined in the SAP. A SAP defined on an interface with encapsulation type Dot1q continues in the dot1d BPDU encapsulation state until a PVST encapsulated BPDU is received 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 — This state specifies that the switch is currently sending proprietary encapsulated BPDUs. PVST BPDUs are only supported on Ethernet interfaces with the encapsulation type set to dot1q. The SAP continues in the PVST BPDU encapsulation state until a dot1d encapsulated BPDU is received, 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.

Configuring VPLS SAPs with per service Split Horizon

Note: Split Horizon group per service is supported only on 7210 SAS-M and 7210 SAS-T devices configured in Network mode.

To configure a VPLS service with a split horizon group, add the **split-horizon-group** parameter when creating the SAP. Traffic arriving on a SAP within a split horizon group will not be copied to other SAPs in the same split horizon group.

The following example displays a 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 SDP Bindings

Note : SDPs are not supported on 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

VPLS provides scaling and operational advantages. A hierarchical configuration eliminates the need for a full mesh of VCs between participating devices. Hierarchy is achieved by enhancing the base VPLS core mesh of VCs with access VCs (spoke) to form two tiers. Spoke SDPs are generally created between Layer 2 switches and placed at the Multi-Tenant Unit (MTU). The PE routers are placed at the service provider's Point of Presence (POP). Signaling and replication overhead on all devices is considerably reduced.

A spoke SDP is treated like the equivalent of a traditional bridge port where flooded traffic received on the spoke SDP is replicated on all other “ports” (other spoke SDPs or SAPs) and not transmitted on the port it was received (unless a split horizon group was defined on the spoke SDP, see section Configuring VPLS Spoke SDPs with Split Horizon on page 353).

A spoke SDP connects a VPLS service between two sites and, in its simplest form, could be a single tunnel LSP. A set of ingress and egress VC labels are exchanged for each VPLS service instance to be transported over this LSP. The PE routers at each end treat this as a virtual spoke connection for the VPLS service in the same way as the PE-MTU connections. This architecture minimizes the signaling overhead and avoids a full mesh of VCs and LSPs between the two metro networks.

A VC-ID can be specified with the SDP-ID. The VC-ID is used instead of a label to identify a virtual circuit. The VC-ID is significant between peer SRs on the same hierarchical level. The value of a VC-ID is conceptually independent from the value of the label or any other datalink specific information of the VC.

Configuring VPLS Spoke SDPs with Split Horizon

Note: Split Horizon group is supported only on 7210 SAS-M and 7210 SAS-T devices configured in Network mode.

To configure spoke SDPs with a split horizon group, add the split-horizon-group parameter when creating the spoke SDP. Traffic arriving on a SAP or spoke SDP within a split horizon group will not be copied to other SAPs or spoke SDPs in the same split horizon group.

The following example displays a 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 50. The tasks below should be performed on both nodes providing the protected VPLS service.

Before configuring a management VPLS, first read VPLS Redundancy on page 292 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 7210 SAS M. Note that the port must be dot1q. The SAP corresponds to the (stacked) VLAN on the 7210 SAS M in which STP is active.
4. Optionally modify STP parameters for load balancing (see Configuring Load Balancing with Management VPLS on page 362).
5. Create an SDP in the m-vpls using the SDP defined in Step 1. Ensure that this SDP runs over a protected LSP.
6. Enable the management VPLS service and verify that it is operationally up.
7. Create a list of VLANs on the port that are to be managed by this management VPLS.
8. Create one or more user VPLS services with SAPs on VLANs in the range defined by Step 6.

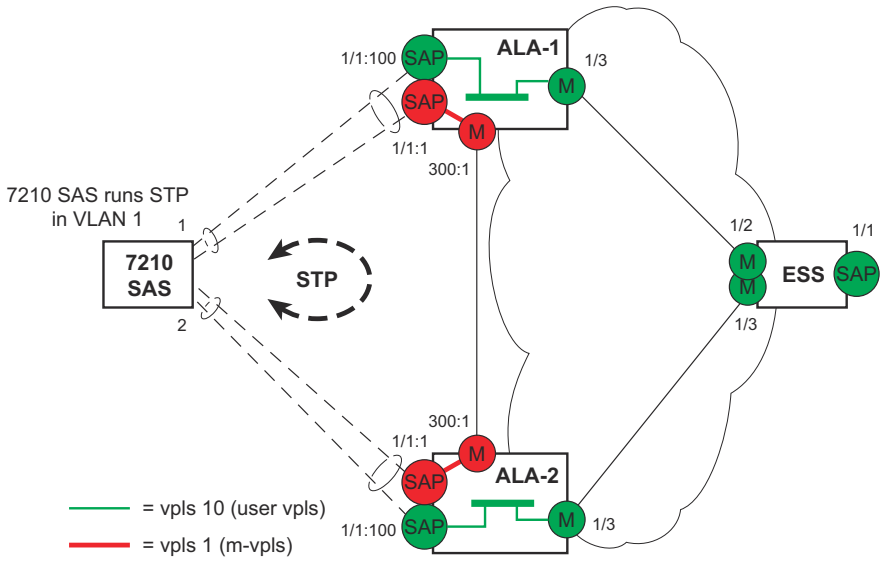

Figure 50: Example Configuration for Protected VPLS SAP

CLI Syntax: config>service# vpls *service-id* [customer *customer-id*] [create] [m-vpls]
description *description-string*
sap *sap-id* create
managed-vlan-list
range *vlan-range*
stp
no shutdown

The following example displays a VPLS configuration:

```
*A:ALA-1>config>service# info
----------------------------------------------
        vpls 2000 customer 6 m-vpls create
            stp
                no shutdown
            exit
            sap 1/1/1:100 create
            exit
            sap 1/1/2:200 create
            exit
            sap 1/1/3:300 create
               managed-vlan-list
                  range 1-50
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-1>config>service#
```

Creating a Management VPLS for Spoke SDP Protection

Note : SDPs are not supported on 7210 SAS devices configured in Access uplink mode. But, Management VPLS can be used for protection of QinQ uplinks. Please refer to the example listed below.

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 51. The tasks below should be performed on all four nodes providing the protected VPLS service.

Before configuring a management VPLS, please first read Configuring a VPLS SAP on page 342 for an introduction to the concept of management VPLS and spoke SDP redundancy.

1. Create an SDP to the local peer node (node ALA-A2 in the example below).
2. Create an SDP to the remote peer node (node ALA-B1 in the example below).
3. Create a management VPLS.
4. Create a spoke SDP in the m-vpls using the SDP defined in Step 1. Ensure that this mesh-spoke SDP runs over a protected LSP (see note below).
5. Enable the management VPLS service and verify that it is operationally up.
6. Create a spoke SDP in the m-vpls using the SDP defined in Step 2.
   Optionally, modify STP parameters for load balancing.
7. Create one or more user VPLS services with spoke SDPs on the tunnel SDP defined by Step 2.

As long as the user spoke SDPs created in step 7are in this same tunnel SDP with the management spoke SDP created in step 6, the management VPLS will protect them.

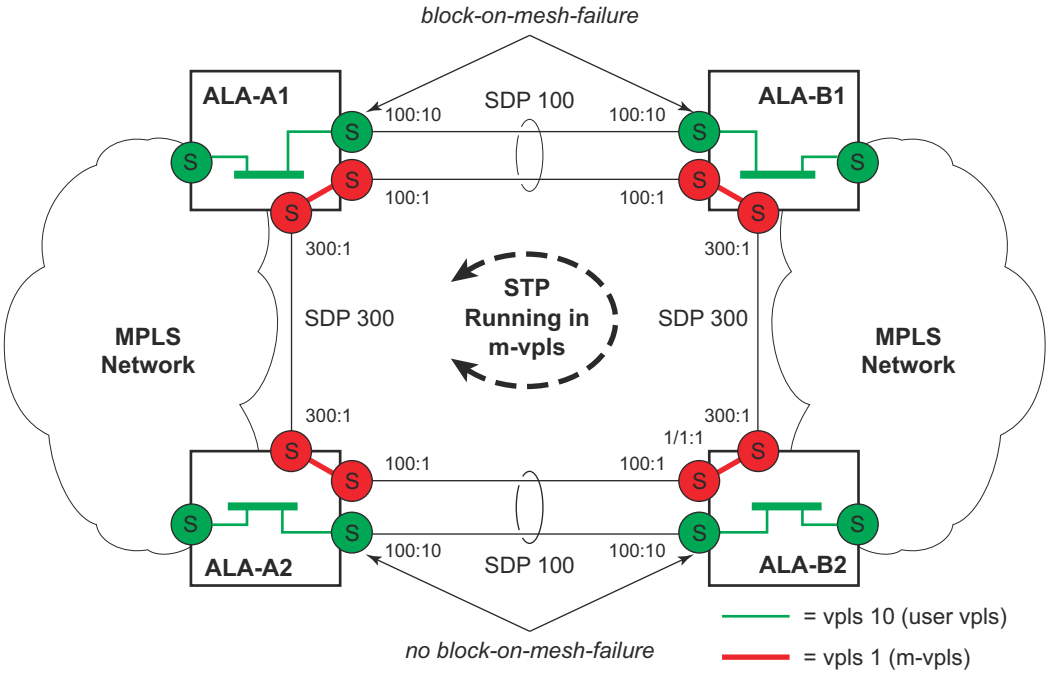

Figure 51: Example Configuration for Protected VPLS Spoke SDP

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
 spoke-sdp sdp-id:vc-id create
 stp
 no shutdown
```

The following example displays a 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
            spoke-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 QinQ spokes (by changing the STP path-cost). When different user VPLS services are associated with either the two management VPLS services, the traffic will be split across the two QinQ spokes. Load balancing can be achieved in SAP protection scenarios.

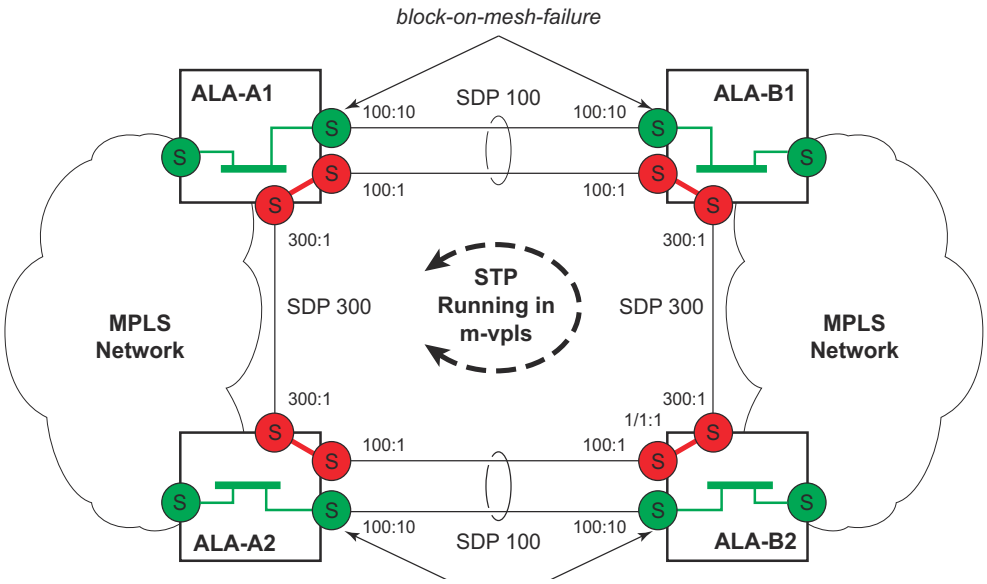

Figure 52: Example Configuration for Load Balancing with Management VPLS

Note: the STP path costs in each peer node should be reversed.

CLI Syntax: config>service# vpls *service-id* [customer *customer-id*] [create] [m-vpls] [svc-sap-type {null-star | dot1q | dot1q-preserve}] [customer-vid vlan-id]

```
            description description-string
            sap sap-id create
               managed-vlan-list
                  range vlan-range
            stp
            no shutdown
```

The following example displays a VPLS configuration:

```
*A:ALA-1>config>service# info
----------------------------------------------
        vpls 100 customer 1 m-vpls svc-sap-type dot1q create
            stp
                no shutdown
            exit
            sap 1/1/2:100.* create
```

```
                managed-vlan-list
                    range 1-10
                exit
                stp
                    path-cost 1
                exit
            exit
            sap 1/1/3:500.* create
                shutdown
                managed-vlan-list
                    range 1-10
                exit
            exit
            no shutdown
        exit
        vpls 200 customer 6 m-vpls svc-sap-type dot1q create
            stp
                no shutdown
            exit
            sap 1/1/2:1000.* create
                managed-vlan-list
                    range 110-200
                exit
            exit
            sap 1/1/3:2000.* create
                managed-vlan-list
                    range 110-200
                exit
                stp
                    path-cost 1
                exit
            exit
            no shutdown
        exit
        vpls 101 customer 1 svc-sap-type dot1q create
            stp
                shutdown
            exit
            sap 1/1/1:100 create
            exit
            sap 1/1/2:1.* create
            exit
            sap 1/1/3:1.* create
            exit
            no shutdown
        exit
        vpls 201 customer 1 svc-sap-type dot1q create
            stp
                shutdown
            exit
            sap 1/1/1:200 create
            exit
            sap 1/1/2:110.* create
            exit
            sap 1/1/3:110.* create
            exit
            no shutdown
        exit
-------------------------------------------

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

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.

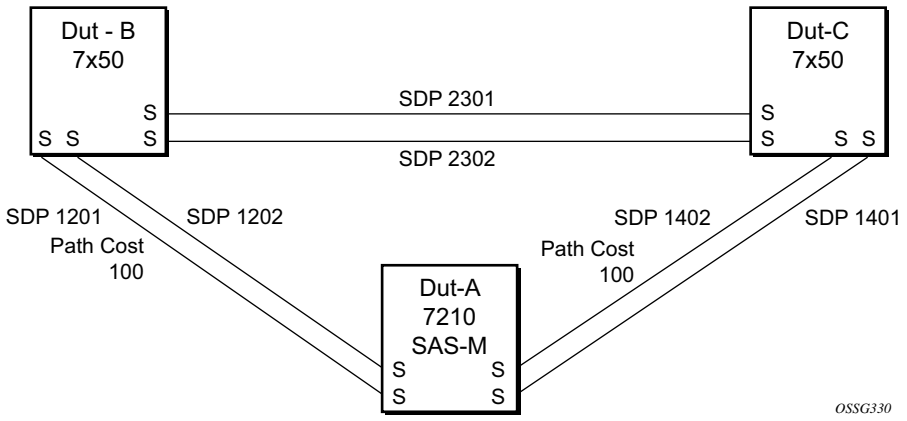


```
mvpls 100                                       mvpls 200
MVPLS M1                                        MVPLSM2
Dut-A Spoke SDP 1201:100 (STP blocked);         Dut-A – Spoke SDP 1202:200; 1402:200 (STP
1401:100                                        blocked)
Dut-B – Spoke SDP 1201:100; 2301:100            Dut-B – Spoke SDP 1202:200; 2302:200
Dut-C – Spoke SDP 1401:100; 2301:100            Dut-C – Spoke SDP 1402:200; 2302:200

uvpls 101                                       uvpls 201
UVPLS U1                                        UVPLS U2
Dut-A – Spoke SDP 1201:101; 1401:101            Dut-A – Spoke SDP 1202:201; 1402:201
Dut-B – Spoke SDP 1201:101; 2301:101            Dut-B – Spoke SDP 1202:201; 2302:201
Dut-C – Spoke SDP 1401:101; 2301:101            Dut-C – Spoke SDP 1402:201; 2302:201
```

Figure 53: 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
    spoke-sdp sdp-id:vc-id create
       stp
          path-cost
    stp
    no shutdown
```

This following output shows example configurations for load balancing across two protected VPLS spoke SDPs:

The configuration on ALA-A (SAS-M) is shown below.

```
# MVPLS 100 configs

*A:ALA-A# configure service vpls 100
*A:ALA-A>config>service>vpls# info
----------------------------------------------
            description "Default tls description for service id 100"
            stp
                no shutdown
            exit
            sap lag-3:100 create
                description "Default sap description for service id 100"
                managed-vlan-list
                    range 101-110
                exit
            exit
            spoke-sdp 1201:100 create
                stp
                    path-cost 100
                exit
            exit
            spoke-sdp 1401:100 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-A>config>service>vpls#


# UVPLS 101 configs

*A:ALA-A>config>service# vpls 101
*A:ALA-A>config>service>vpls# info
----------------------------------------------
            description "Default tls description for service id 101"
            sap lag-3:101 create
                description "Default sap description for service id 101"
            exit
            spoke-sdp 1201:101 create
```

```
            exit
            spoke-sdp 1401:101 create
            exit
            no shutdown
-----------------------------------------------
*A:ALA-A>config>service>vpls#


# MVPLS 200 configs

*A:ALA-A# configure service vpls 200
*A:ALA-A>config>service>vpls# info
-----------------------------------------------
            description "Default tls description for service id 200"
            stp
                no shutdown
            exit
            sap lag-3:200 create
                description "Default sap description for service id 200"
                managed-vlan-list
                    range 201-210
                exit
            exit
            spoke-sdp 1202:200 create
            exit
            spoke-sdp 1402:200 create
                stp
                    path-cost 100
                exit
            exit
            no shutdown
-----------------------------------------------
*A:ALA-A>config>service>vpls#


# UVPLS 201 configs

*A:ALA-A>config>service# vpls 201
*A:ALA-A>config>service>vpls# info
-----------------------------------------------

            description "Default tls description for service id 201"
            sap lag-3:201 create
                description "Default sap description for service id 201"
            exit
            spoke-sdp 1202:201 create
            exit
            spoke-sdp 1402:201 create
            exit
            no shutdown
-----------------------------------------------
*A:ALA-A>config>service>vpls# exit all
```

The configuration on ALA-B (7x50), the top left node is shown below. It is configured such that it becomes the root bridge for MVPLS 100 and MVPLS 200.

```
# MVPLS 100 configs

*A:ALA-B# configure service vpls 100
*A:ALA-B>config>service>vpls# info
----------------------------------------------
            description "Default tls description for service id 100"
            stp
                priority 0
                no shutdown
            exit
            spoke-sdp 1201:100 create
            exit
            spoke-sdp 2301:100 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-B>config>service>vpls#


# UVPLS 101 configs

*A:ALA-B>config>service# vpls 101
*A:ALA-B>config>service>vpls# info
----------------------------------------------

            description "Default tls description for service id 101"
            spoke-sdp 1201:101 create
            exit
            spoke-sdp 2301:101 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-B>config>service>vpls#


# MVPLS 200 configs

*A:ALA-B# configure service vpls 200
*A:ALA-B>config>service>vpls# info
----------------------------------------------
            description "Default tls description for service id 200"
            stp
                priority 0
                no shutdown
            exit
            spoke-sdp 1202:200 create
            exit
            spoke-sdp 2302:200 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-B>config>service>vpls#
```

```
# UVPLS 201 configs

*A:ALA-B>config>service# vpls 201
*A:ALA-B>config>service>vpls# info
----------------------------------------------

            description "Default tls description for service id 201"
            spoke-sdp 1202:201 create
            exit
            spoke-sdp 2302:201 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-B>config>service>vpls#
```

The configuration on ALA-C (7210), the top right node is shown below.

```
# MVPLS 100 configs

*A:ALA-C# configure service vpls 100
*A:ALA-C>config>service>vpls# info
----------------------------------------------

            description "Default tls description for service id 100"
            stp
                priority 4096
                no shutdown
            exit
            spoke-sdp 1401:100 create
            exit
            spoke-sdp 2301:100 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-C>config>service>vpls#


# UVPLS 101 configs

*A:ALA-C>config>service# vpls 101
*A:ALA-C>config>service>vpls# info
----------------------------------------------

            description "Default tls description for service id 101"
            spoke-sdp 1401:101 create
            exit
            spoke-sdp 2301:101 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-C>config>service>vpls#
```

```
# MVPLS 200 configs

*A:ALA-C# configure service vpls 200
*A:ALA-C>config>service>vpls# info
----------------------------------------------
            description "Default tls description for service id 200"
            stp
                priority 4096
                no shutdown
            exit
            spoke-sdp 1402:200 create
            exit
            spoke-sdp 2302:200 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-C>config>service>vpls#


# UVPLS 201 configs

*A:ALA-C>config>service# vpls 201
*A:ALA-C>config>service>vpls# info
----------------------------------------------

            description "Default tls description for service id 201"
            spoke-sdp 1402:201 create
            exit
            spoke-sdp 2302:201 create
            exit
            no shutdown
----------------------------------------------
*A:ALA-C>config>service>vpls#
```

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`

Configuring Load Balancing with Management VPLS

With the concept of management VPLS, it is possible to load balance the user VPLS services across the two protecting nodes. This is done by creating two management VPLS instances, where both instances have different active QinQ spokes (by changing the STP path-cost). When different user VPLS services are associated with either the two management VPLS services, the traffic will be split across the two QinQ spokes. Load balancing can be achieved in SAP protection scenarios.

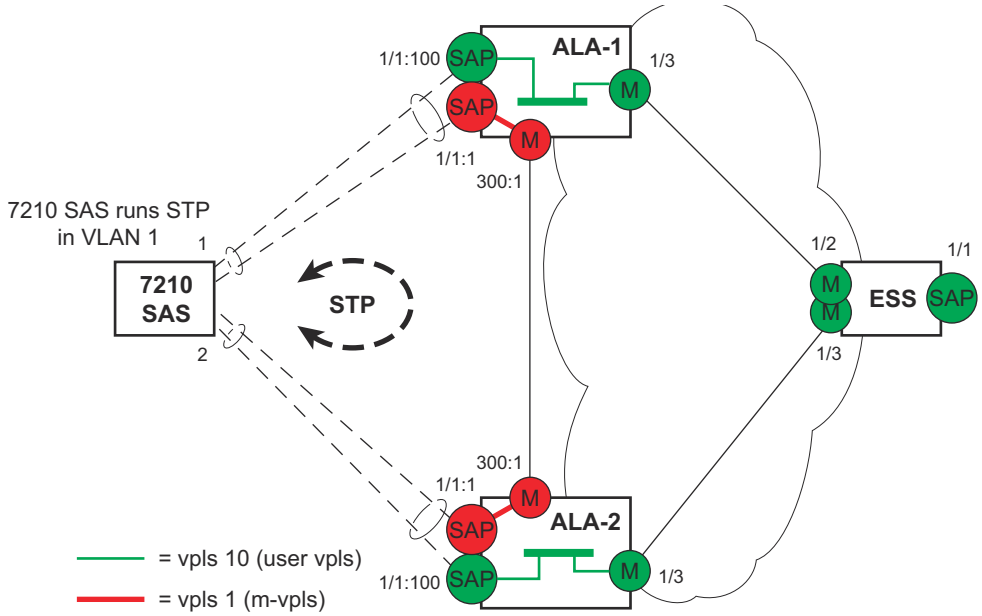

Figure 54: Example Configuration for Load Balancing with Management VPLS

Note: the STP path costs in each peer node should be reversed.

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

```
            description description-string
            sap sap-id create
               managed-vlan-list
                  range vlan-range
            stp
            no shutdown
```

The following example displays a VPLS configuration:

```
*A:ALA-1>config>service# info
----------------------------------------------
        vpls 100 customer 1 m-vpls svc-sap-type any create
            stp
                no shutdown
            exit
            sap 1/1/2:100.* create
                managed-vlan-list
```

```
                        range 1-10
                    exit
                    stp
                        path-cost 1
                    exit
                exit
                sap 1/1/3:500.* create
                    shutdown
                    managed-vlan-list
                        range 1-10
                    exit
                exit
                no shutdown
            exit
            vpls 200 customer 6 m-vpls svc-sap-type any create
                stp
                    no shutdown
                exit
                sap 1/1/2:1000.* create
                    managed-vlan-list
                        range 110-200
                    exit
                exit
                sap 1/1/3:2000.* create
                    managed-vlan-list
                        range 110-200
                    exit
                    stp
                        path-cost 1
                    exit
                exit
                no shutdown
            exit
            vpls 101 customer 1 svc-sap-type any create
                stp
                    shutdown
                exit
                sap 1/1/1:100 create
                exit
                sap 1/1/2:1.* create
                exit
                sap 1/1/3:1.* create
                exit
                no shutdown
            exit
            vpls 201 customer 1 svc-sap-type any create
                stp
                    shutdown
                exit
                sap 1/1/1:200 create
                exit
                sap 1/1/2:110.* create
                exit
                sap 1/1/3:110.* create
                exit
                no shutdown
            exit
        ----------------------------------------------

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

Configuring BGP Auto-Discovery

This section provides important information to explain the different configuration options used to populate the required BGP AD and generate the LDP generalized pseudowire-ID FEC fields. There are a large number of configuration options that are available with the this feature. Not all these configurations option are required to start using BGP AD. At the end of this section, it will be apparent that a very simple configuration will automatically generate the required values used by BGP and LDP. In most cases, deployments will provide full mesh connectivity between all nodes across a VPLS instance. However, capabilities are available to influence the topology and build hierarchies or hub and spoke models.

Configuration Steps

Using Figure 55, assume PE6 was previously configured with VPLS 100 as indicated by the configurations lines in the upper right. The BGP AD process will commence after PE134 is configured with the VPLS 100 instance as shown in the upper left. This shows a very basic and simple BGP AD configuration. The minimum requirement for enabling BGP AD on a VPLS instance is configuring the VPLS-ID and point to a pseudowire template.

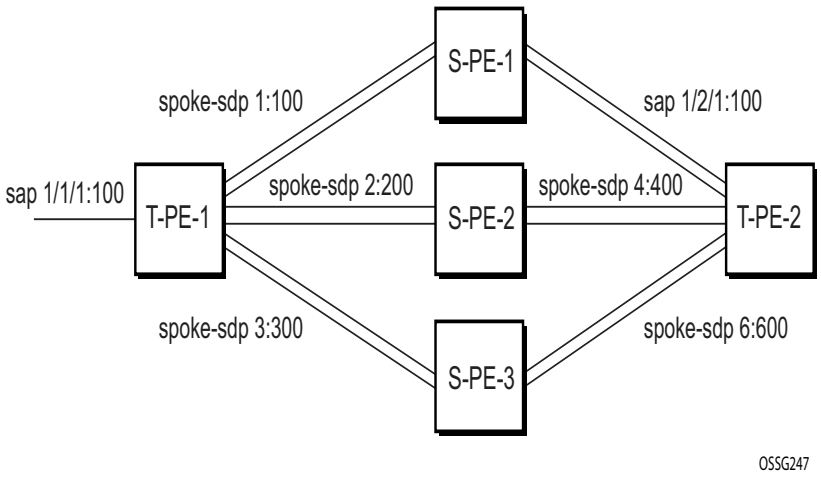

Figure 55: BGP AD Configuration Example

In many cases, VPLS connectivity is based on a pseudowire mesh. To reduce the configuration requirement, the BGP values can be automatically generated using the VPLS-ID and the MPLS router-ID. By default, the lower six bytes of the VPLS-ID are used to generate the RD and the RT values. The VSI-ID value is generated from the MPLS router-ID. All of these parameters are configurable and can be coded to suit requirements and build different topologies

```
PE134>config>service>vpls>bgp-ad#
 [no] pw-template-bi* - Configure pw-template bind policy
 [no] route-target    - Configure route target
 [no] shutdown        - Administratively enable/disable BGP auto-discovery
      vpls-id         - Configure VPLS-ID
 [no] vsi-export      - VSI export route policies
      vsi-id          + Configure VSI-id
 [no] vsi-import      - VSI import route policies
```

Figure 56: BGP-AD CLI Command Tree

A helpful command displays the service information, the BGP parameters and the SDP bindings in use. When the discovery process is completed successfully each endpoint will have an entry for the service.

```
PE134>#  show service l2-route-table
```

When only one of the endpoints has an entry for the service in the l2-routing-table, it is most likely a problem with the RT values used for import and export. This would most likely happen when different import and export RT values are configured using a router policy or the route-target command.

Service specific commands continue to be available to display service specific information, including status.

```
PERs6# show service sdp-using
```

BGP AD advertises the VPLS-ID in the extended community attribute, VSI-ID in the NLRI and the local PE ID in the BGP next hop. At the receiving PE, the VPLS-ID is compared against locally provisioned information to determine whether the two PEs share a common VPLS. If it is found that they do, the BGP information is used in the signaling phase.

Configuring AS Pseudo-wire in VPLS

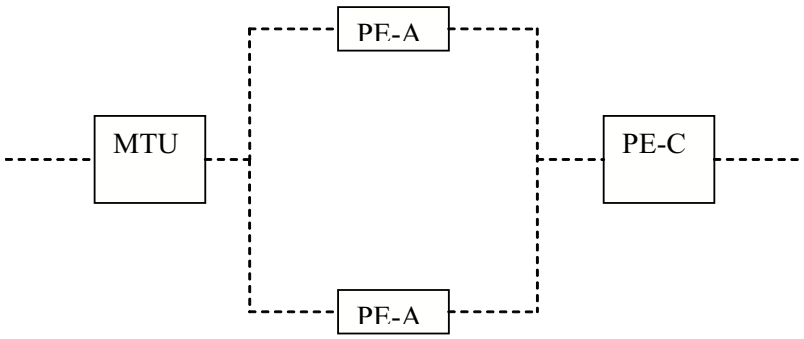

Figure 57: Sample Topology-AS Pseudo-wire in VPLS

In Figure 57, Pseudo-wire is configured on MTU. A sample configuration on the MTU is listed below:

```
*A:MTU>config>service>vpls>endpoint# back
*A:MTU>config>service>vpls# info
----------------------------------------------
            send-flush-on-failure
            stp
                shutdown
            exit
            endpoint "vpls1" create
                description "vpls1_endpoint"
                revert-time 60
                ignore-standby-signaling
                no suppress-standby-signaling
                block-on-mesh-failure
            exit
            sap 1/1/3 create
            exit
            spoke-sdp 301:1 endpoint "vpls1" create
                stp
                    shutdown
                exit
                block-on-mesh-failure
            exit
            spoke-sdp 302:1 endpoint "vpls1" create
                stp
                    shutdown
                exit
                block-on-mesh-failure
            exit
            no shutdown
----------------------------------------------
*A:MTU>config>service>vpls#
```

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 and then enter the new information.

The following displays a modified VPLS configuration.

```
*A:ALA-1>config>service>vpls# info
----------------------------------------------
            description "This is a different description."
            disable-learning
            disable-aging
            discard-unknown
            local-age 500
            stp
                shutdown
            exit
            sap 1/1/5:22 create
                description "VPLS SAP"
            exit
            exit
            no shutdown
----------------------------------------------
*A:ALA-1>config>service>vpls#
```

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
```

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] spoke-sdp sdp-id
        [no] sap sap-id
           shutdown
```

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 (not applicable for 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode) 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] spoke-sdp sdp-id
          shutdown
       sap sap-id
       no sap sap-id
          shutdown
```

Disabling a VPLS Service

You can shut down a VPLS service without deleting the service parameters.

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

Example:
```
config>service# vpls 1
config>service>vpls# shutdown
config>service>vpls# exit
```

Re-Enabling a VPLS Service

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

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

Example:
```
config>service# vpls 1
config>service>vpls# no shutdown
config>service>vpls# exit
```

VPLS Services Command Reference

Command Hierarchies

VPLS Service Configuration Commands

Global Commands

```
config
    — service
        — vpls service-id [customer customer-id] [create [r-vpls] (for 7210 SAS-M and 7210 SAS-T in Network mode)
        — vpls service-id [customer customer-id] [create] [ vpn vpn-id] [m-vpls] <service-id> [svc-sap-type {null-star|dot1q-preserve|dot1q-range|any}] [customer-vid <vlan-id>] (for 7210 SAS-M and 7210 SAS-T in Access uplink mode)
        — no vpls service-id
            — bgp
                — pw-template-binding policy-id [split-horizon-group group-name] [import-rt {ext-community...(up to 5 max)}]
                — no pw-template-binding policy-id
                — route-distinguisher [ip-addr:comm-val | as-number:ext-comm-val]
                — no route-distinguisher
                — route-target {ext-community | {[export ext-community] [import ext-community]}}
                — no route-target
                — vsi-export policy-name [policy-name...(up to 5 max)]
                — no vsi-export
                — vsi-import policy-name [policy-name...(up to 5 max)]
                — no vsi-import
            — [no] bgp-ad
                — [no] shutdown
                — vpls-id vpls-id
                — vsi-id
                    — prefix low-order-vsi-id
                    — no prefix
            — description description-string
            — no description
            — [no] disable-aging
            — [no] disable-learning
            — [no] discard-unknown
            — endpoint endpoint-name [create]
            — no endpoint
                — block-on-mesh-failure
                — [no] block-on-mesh-failure
                — description description-string
                — no description
                — [no] ignore-standby-signaling
                — [no] mac-pinning
                — max-nbr-mac-addr table-size
                — no max-nbr-mac-addr
                — revert-time revert-time | infinite
                — no revert-time
                — static-mac ieee-address [create]
                — no static-mac
                — [no] suppress-standby-signaling
            — eth-cfm
                — [no] mep mep-id domain md-index association ma-index [direction {up|down}]
                — [no] mep mep-id domain md-index association ma-index
```

- [**no**] **ccm-enable**
- **ccm-ltm-priority** *priority*
- **no ccm-ltm-priority**
- [**no**] **description**
- [**no**] **eth-test-enable**
  - [**no**] **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
- **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}
- **mac-address** *mac-address*
- **no mac-address**
- **one-way-delay-threshold** *seconds*
- [**no**] **shutdown**

tunnel-fault [**accept** | **ignore**]

[**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*]

igmp-snooping

- **mvr**
  - **description** *description-string*
  - **no description**
  - **group-policy** *policy-name*
  - **no group-policy**
  - [**no**] **shutdown**
- **query-interval** *seconds*
- **no query-interval**
- **query-src-ip** *seconds*
- **no query-src-ip**
- **report-src-ip** *ip-address*
- **no report-src-ip**
- **robust-count** *robust-count*
- **no robust-count**
- [**no**] **shutdown**

[**no**] **interface** *ip-int-name* [**create**] (for 7210 SAS-M and 7210 SAS-T in access uplink mode)

- **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
— 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] propagate-mac-flush
— remote-age aging-timer
— no remote-age
— [no] send-flush-on-failure
— service-mtu octets (for 7210 SAS-M and 7210 SAS-T in Network mode)
— no service-mtu
— no service-mtu-check (for 7210 SAS-M and 7210 SAS-T in Network mode)
— [no] shutdown
— split-horizon-group group-name [create]
            — description description-string
            — no description
— stp
            — forward-delay forward-delay
            — no forward-delay
            — hello-time hello-time
            — no hello-time
            — hold-count BDPU tx hold count
            — no hold-count
            — max-age max-age
            — no max-age
            — mode {rstp | comp-dot1w | dot1w | mstp | pmstp}
            — no mode
            — [no] mst-instance mst-inst-number
                        — mst-port-priority bridge-priority
                        — no mst-port-priority
                        — [no] vlan-range vlan-range
            — mst-max-hops hops-count
            — no mst-max-hops
            — mst-name region-name
            — no mst-name
            — mst-revision revision-number
            — no mst-revision
            — priority bridge-priority
            — no priority
            — [no] shutdown
```

SAP Commands

config
- **service**
  - **vpls** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**m-vpls**] (for 7210 SAS-M and 7210 SAS-T in network mode)
  - **vpls** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] *service-id* [**create**] [**vpn** *vpn-id*] [**m-vpls**] [**svc-sap-type** {**null-star** | **dot1q** | **dot1q-preserve**}] [**customer-vid** *vlan-id*](for 7210 SAS-M and 7210 SAS-T in access uplink mode)
  - **no vpls** *service-id*
    - **sap** *sap-id* [**split-horizon-group** *group-name*] [g8032-shg-enable][**eth-ring** *ring-index*] [**create**]
    - **no sap** *sap-id*
      - **accounting-policy** *acct-policy-id*
      - **no accounting-policy**
      - **bpdu-translation** {**auto** | **pvst** | **stp**}
      - **no bpdu-translation**
      - [**no**] **collect-stats**
      - **description** *description-string*
      - **no description**
      - **dhcp**
        - **description** *description-string*
        - **no description**
        - [**no**] **option**
          - **action** [*dhcp-action*]
          - **no action**
          - **[no] circuit-id** [**ascii-tuple** | **vlan-ascii-tuple**]
          - [**no**] **remote-id** [**mac** | **string** *string*]
          - [**no**] **vendor-specific-option**
            - [**no**] **client-mac-address**
            - [**no**] **sap-id**
            - [**no**] **service-id**
            - **string** *text*
            - **no string**
            - [**no**] **system-id**
        - [**no**] **shutdown**
        - [**no**] **snoop**
      - [**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*] [**ipv6** *ipv6 -filter-id*] [**mac** *mac-filter-id*]
      - **eth-cfm**
        - **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** **{**up | down**}**] **primary-vlan-enable** [**vlan** vlan-id]
        - **no mep** *mep-id* **domain** *md-index* **association** *ma-index*
          - [**no**] **ais-enable**
            - **client-meg-level** [*level* [*level...*]]
            - **no client-meg-level**
            - [**no**] **description**
            - **interval** {**1**| **60**}

- **no interval**
- **priority** *priority-value*
- **no priority**

- [**no**] **ccm-enable**
- **ccm-ltm-priority** *priority*
- **no ccm-ltm-priority**
- **description** *description-string*
- **no description**
- [**no**] **eth-test-enable**
  - **bit-error-threshold** *bit-errors*
  - **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
  - **no test-pattern**
- **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}
- **mac-address** *mac-address*
- **no mac-address**
- **one-way-delay-threshold** *seconds*
- [**no**] **shutdown**

- **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**
  - [**no**] **mrouter-port**
  - **mvr**
    - **from-vpls** *service-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* (applicable only in access-uplink mode)
        - [**no**] **starg**
  - **version** *version*
  - **no version**
- **ingress**
  - **aggregate-meter-rate** *rate-in-kbps* [**burst** *burst-in-kbits*]
  - **no aggregate-meter-rate**
  - **filter ip** *ip-filter-id*
  - **filter [ipv6** *ipv6-filter-id***]**
  - **filter mac** *mac-filter-id*
  - **no filter** [**ip** *ip-filter-id*] **[ipv6** *ipv6-filter-id***]** [**mac** *mac-filter-id*]
  - **meter-override**
  - **no meter-override**

```
                        — meter  meter-id [create]
                        — no meter meter-id
                            — adaptation-rule [pir adaptation-rule] [cir adap-
                              tation-rule]
                            — cbs size-in-kbytes
                            — no cbs
                            — mbs size-in-kbits
                            — no mbs
                            — mbs mode
                            — no mode
                            — no mode
                            — rate cir cir-rate [pir pir-rate]
                    — qos policy-id
                    — no qos
            — l2pt-termination [cdp] [dtp] [pagp] [stp] [udld] [vtp]
            — 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
            — [no] shutdown
            — statistics
                    — ingress
                        — counter-mode {in-out-profile-count|forward-drop-
                          count}
                        — drop-count-extra-vlan-tag-pkts
                        — no drop-count-extra-vlan-tag-pkts
            — stp
                    — [no] auto-edge
                    — [no] edge-port
                    — link-type {pt-pt | shared}
                    — no link-type [pt-pt | shared]
                    — mst-instance mst-inst-number
                        — mst-path-cost inst-path-cost
                        — no mst-path-cost
                        — mst-port-priority stp-priority
                        — no mst-port-priority
                    — path-cost sap-path-cost
                    — no path-cost
                    — [no] port-num virtual-port-number
                    — priority stp-priority
                    — no priority
                    — no root-guard
                    — root-guard
                    — [no] shutdown
            — tod-suite tod-suite-name
            — no tod-suite
```

Mesh SDP Commands

Note: Mesh SDP commands are not supported by 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode. It is supported only in 7210 SAS-M and 7210 SAS-T network mode.

config
- **service**
  - [**no**] **vpls** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**mvpls**] [create] [vpn vpn-id] [m-vpls] (for 7210 SAS-M and 7210 SAS-T in Network mode)
    - **mesh-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}]
    - **no mesh-sdp** *sdp-id*[:*vc-id*]
      - **accounting-policy** *acct-policy-id*
      - **no accounting-policy**
      - [**no**] **collect-stats**
      - [**no**] **control-word**
      - **description** *description-string*
      - **no description**
      - **egress**
        - **no vc-label** [*egress-vc-label*]
      - **eth-cfm**
        - **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {**up**}{**down**}]
        - **no mep** *mep-id* **domain** *md-index* **association** *ma-index*
          - [**no**] **ais-enable**
            - **client-meg-level** [[*level* [*level*...]]
            - **no client-meg-level**
            - **interval** {**1** | **60**}
            - **no interval**
            - **priority** *priority-value*
            - **no priority**
          - [**no**] **ccm-enable**
          - **ccm-ltm-priority** ***priority***
          - **no ccm-ltm-priority**
          - [**no**] **description** *description-string*
          - [**no**] **eth-test-enable**
            - **bit-error-threshold** *bit-errors*
            - **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
            - **no test-pattern**
          - **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}
          - **mac-address** ***mac-address***
          - **no mac-address**
          - **one-way-delay-threshold** *seconds*
          - [**no**] **shutdown**
      - [**no**] **force-vlan-vc-forwarding**
      - **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**
        - [**no**] **mrouter-port**
        - **query-interval** *interval*
        - **no query-interval**

```
            — query-response-interval interval
            — no query-response-interval
            — robust-count count
            — no robust-count
            — [no] send-queries
            — static
                — [no] group grp-ip-address
                — [no] starg
            — version version
            — no version
— ingress
            — vc-label egress-vc-label
— [no] mac-pinning
— [no] static-mac ieee-address
— [no] static-mac ieee-address [create][no] shutdown
— statistics
            — ingress[no] drop-count-extra-vlan-tag-pkts
— vlan-vc-tag 0..4094
— no vlan-vc-tag [0..4094]
```

Spoke SDP Commands

Note: Spoke SDP commands are not supported by 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode. It is supported only in network mode.

config
- **service**
  - [**no**] **vpls** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**mvpls**] (for 7210 SAS-M and 7210 SAS-T in Network mode)
    - **spoke-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] [**split-horizon-group** *group-name*] **endpoint** *endpoint-name*
    - **no spoke-sdp** *sdp-id*[:*vc-id*]
      - **accounting-policy** *acct-policy-id*
      - **no accounting-policy**
      - [**no**] **block-on-mesh-failure**
      - **bpdu-translation** {**auto** | **pvst** | **stp**}
      - **no bpdu-translation**
      - [**no**] **collect-stats**
      - [**no**] **control-word**
      - **description** *description-string*
      - **no description**
      - [**no**] **disable-aging**
      - [**no**] **disable-learning**
      - [**no**] **discard-unknown-source**
      - **eth-cfm**
        - **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {**up**}{**down**}]
          - **no mep** *mep-id* **domain** *md-index* **association** *ma-index*[**no**] **ais-enable**
            - **client-meg-level** [[*level* [*level...*]]
            - **no client-meg-level**
            - **interval** {**1** | **60**}
            - **no interval**
            - **priority** *priority-value*
            - **no priority**
          - [**no**] **ccm-enable**
          - **ccm-ltm-priority** ***priority***
          - **no ccm-ltm-priority**
          - [**no**] **description** *description string*[**no**] **eth-test-enable**
            - **bit-error-threshold** *bit-errors*
            - **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
          - **no test-patternlow-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}
          - **mac-address** *mac-address*
          - **no mac-address****one-way-delay-threshold** *seconds*
          - [**no**] **shutdown**
      - **egress**
        - **vc-label** *egress-vc-label*
        - **no vc-label** [*egress-vc-label*]
      - [**no**] **force-vlan-vc-forwarding**
      - **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
            — [no] mrouter-port
            — query-interval interval
            — no query-interval
            — query-response-interval interval
            — no query-response-interval
            — robust-count count
            — no robust-count
            — [no] send-queries
            — static
                  — [no] group group-address
                        — [no] starg
            — version version
            — no version
      — [no] ignore-standby-signaling
      — ingress
            — vc-label egress-vc-label
            — no vc-label [egress-vc-label]
      — [no] l2pt-termination
      — limit-mac-move [blockable | non-blockable]
      — no limit-mac-move
      — [no] mac-pinning
      — max-nbr-mac-addr table-size
      — no max-nbr-mac-addr
      — precedence precedence-value | primary
      — no precedence
      — [no] shutdown
      — [no] static-mac ieee-address
      — statistics
            — ingress
            — [no] drop-count-extra-vlan-tag-pkts
      — stp
            — [no] auto-edge
            — [no] edge-port
            — link-type {pt-pt | shared}
            — no link-type [pt-pt | shared]
            — path-cost sap-path-cost
            — no path-cost
            — [no] port-num virtual-port-number
            — priority stp-priority
            — no priorityno root-guard
            — root-guard
            — [no] shutdown
      — vlan-vc-tag 0..4094
      — no vlan-vc-tag [0..4094]
```

Routed VPLS Commands applicable only to 7210 SAS- M and T

NOTE: The command “allow-ip-int-binding” is applicable for 7210 SAS-M and 7210 SAS-T devices in Access-Uplink mode.

config
- **service**
  - **vpls** *service-id* **[customer** *customer-id***] [vpn** *vpn-id***] [m-vpls] [create]**
    - **service-name** *service-name*
    - **no service-name**
      - **[no] allow-ip-int-binding**

Routed VPLS Commands (Supported on 7210 SAS-M (Network mode) and 7210 SAS-X)

config
- **service**
  - **vpls** *service-id* **[customer** *customer-id***] [create] [vpn** *vpn-id***] [m-vpls] [svc-sap-type {null-star|dot1q-preserve|any}] [customer-vid** *vlan-id***] [b-vpls|i-vpls|r-vpls]**
    - **service-name** *service-name*
    - **no service-name**
      - **[no] allow-ip-int-binding**

Show Commands

Note : SDP commands are not supported by 7210 SAS devices configured in Access uplink mode.

show
- **service**
  - **egress-label** *egress-label1* [*egress-label2*]
  - **fdb-info**
  - **fdb-mac** *ieee-address* [**expiry**]
  - **id** *service-id*
    - **all**
    - **base** [**msap**] [**bfd**]
    - **dhcp**
      - **statistics** [**sap** *sap-id*] [**interface** *interface-name*]
      - **summary [interface** *interface-name* **| saps]**
    - **endpoint** [*endpoint-name*]
    - **fdb** [**sap** *sap-id*] [**expiry**]] | [**mac** *ieee-address* [**expiry**]] | [**detail**] [**expiry**]
    - **igmp-snooping**
      - **all**
      - **base**
      - **mvr**
      - **mrouters** [**detail**]
      - **port-db sap** *sap-id* [**detail**]
      - **port-db sap** *sap-id* **group** *grp-address*
      - **port-db sdp** *sdp-id*:*vc-id* [**detail**]
      - **port-db sdp** *sdp-id*:*vc-id* **group** *grp-address*
      - **proxy-db** [**detail**]
      - **proxy-db** [**group** *grp-ip-address*]
      - **querier**
      - **static** [**sap** *sap-id*]
      - **statistics**[**sap** *sap-id* | **sdp** *sdp-ic:vc-id*]
    - **labels**
    - **l2pt disabled**
    - **l2pt** [**detail**]
    - **mac-move**
    - **mfib** [**brief**]
    - **mfib** [**group** *grp-address* | **mstp-configuration**]
    - **sap** [*sap-id* [**detail**]]
    - **sdp** [*sdp-id* | **far-end** *ip-addr*] [**detail**]
    - **split-horizon-group** [*group-name*]
    - **stp** [**detail**]
  - **ingress-label** *start-label* [*end-label*]
  - **sap-using** [**sap** *sap-id*]
  - **sap-using** [**ingress** | **egress**] **filter** ***f****ilter-id*
  - **sap-using** [**ingress** | **egress**] **qos-policy** *qos-policy-id*
  - **sap-using** [**ingress** | **egress**]
  - **sdp** [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]
  - **sdp-using** [*sdp-id*[:*vc-id*] | **far-end** *ip-address*]
  - **service-using** [**vpls**]

Clear Commands

Note : SDP commands are not supported by 7210 SAS devices configured in Access uplink mode.

clear
- **service**
  - **id** *service-id*
    - **fdb** {**all** | **mac** *ieee-address* | **sap** *sap-id* | **mesh-sdp** *sdp-id*[**:***vc-id*] | **spoke-sdp** *sdp-id***:***vc-id*}
    - **igmp-snooping**
      - **port-db sap *sap-id*** [**group** *grp-address*]
      - **querier**
    - **statistics** [**all** | **sap** *sap-id* | **sdp** *sdp-id*:*vc-id*]
    - **mesh-sdp** *sdp-id*[**:***vc-id*] **ingress-vc-label**
      - **spoke-sdp** *sdp-id***:***vc-id* **ingress-vc-label**
    - **spoke-sdp** *sdp-id[:vc-id]*
    - **stp**
      - **detected-protocols** [**all** | **sap** *sap-id*]
  - **statistics**
    - **id** *service-id*
      - **cem (applicable only for 7210 SAS-M and 7210 SAS-T in Network mode)**
      - **counters**
      - **mesh-sdp sdp-id**[:*vc-id*] {**all** | **counters** | **stp** }
      - **spoke-sdp** *sdp-id*[:*vc-id*] {**all** | **counters** | **stp** | **l2pt**}
      - **stp**
    - **sap** *sap-id* {**all** | **counters** | **stp** }
    - **sdp** *sap-id* {**keep-alive**}

Debug Commands

debug
— service
— id *service-id*

VPLS Service Configuration Commands

Generic Commands

shutdown

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

Context config>service>vpls
config>service>vpls>snooping
config>service>vpls>igmp-snooping
config>service>vpls>sap
config>service>vpls>sap>stp
config>service>vpls>stp
config>service>vpls>spoke-sdp>stp
config>service>vpls>bgp-ad

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.

description

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

Context config>service>vpls
config>service>vpls>split-horizon-group
config>service>vpls>igmp-snooping>mvr
config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>pw-template>split-horizon-group

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.

DHCP Commands

dhcp

Syntax **dhcp**

Context config>service>vpls>sap
config>service>ies>sap
config>service>vprn>sap

Description This command enables the context to configure DHCP parameters.

option

Syntax [**no**] **option**

Context config>service>vpls>sap>dhcp
config>service>ies>sap>dhcp
config>service>vprn>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*] {***replace*** | ***drop*** | ***keep***}
no action

Context config>service>vpls>sap>dhcp>option
config>service>ies>sap>dhcp>option
config>service>vprn>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.

gi-address

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

Context config>service>ies>if>dhcp
config>service>vprn>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.

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.

circuit-id

Syntax **circuit-id** [**ascii-tuple** | **ifindex** | **sap-id** | **vlan-ascii-tuple**]
no circuit-id

Context config>service>vpls>sap>dhcp>option
config>service>ies>if>dhcp>option
config>service>vprn>if>dhcp>option

When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by "|".

In order to send a tuple in the circuit ID, the **action replace** command must be configured in the same context.

If disabled, the **circuit-id** sub-option 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 — 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.

vlan-ascii-tuple — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q-encapsulated ports only. 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
config>service>ies>sap>dhcp>option
config>service>vprn>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

config>service>vprn>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
config>service>vpls>sap>dhcp>option>vendor
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
config>service>ies>sap>dhcp>option>vendor
config>service>vprn>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
config>service>ies>sap>dhcp>option>vendor
config>service>vprn>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
config>service>ies>sap>dhcp>option>vendor
config>service>vprn>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
config>service>ies>sap>dhcp>option>vendor
config>service>vprn>sap>dhcp>option>vendor

Description This command specifies whether the system-id is encoded in the vendor specific sub-option of Option 82.

server

Syntax **server** *server1* [*server2*...(up to 8 max)]

Context config>service>ies>if>dhcp
config>service>vprn>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.

Default no server

Parameters *server* — Specify the DHCP server IP address.

trusted

Syntax [**no**] **trusted**

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

Description This command enables relaying of untrusted packets.

The **no** form of this command disables the relay.

Default not enabled

snoop

Syntax [**no**] **snoop**

config>service>vpls>sap>dhcp **Context**

Description This command enables DHCP snooping of DHCP messages on the SAP. Enabling DHCP snooping on VPLS interfaces (SAPs) is required where DHCP messages where Option 82 information is to be inserted. This includes interfaces that are in the path to receive messages from either DHCP servers or from subscribers.

Use the **no** form of the command to disable DHCP snooping on the specified VPLS SAP.

Note: The 7210 SAS-X and 7210 SAS-M and 7210 SAS-T (network mode) does not support DHCP snooping for SDP.

Default no snoop

VPLS Service Commands

vpls

Syntax **vpls** *service-id* [**customer** *customer-id*] [**create**] *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**m-vpls] [r-vpls**] (for 7210 SAS-M and 7210 SAS-T in network mode)
vpls *service-id* [**customer** *customer-id*] [**create**][**vpn** *vpn-id*] [*m-vpls*] *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [m-vpls] [**svc-sap-type** {**null-star**|**dot1q-preserve|dot1q-range**|**any**}] [**customer-vid** *vlan-id*] (for 7210 SAS-M and 7210 SAS-T in access uplink mode)
vpls *service-id* **customer** *customer-id* **vpn** *vpn-id* **[m-vpls] [bvpls | i-vpls] [create]**
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 defined within the service ID have been shutdown and deleted, and the service has been shutdown.

Parameters *any* — Allows any SAP type. When svc-sap-type is set to **any**, for a NULL SAP, the system processes and forwards only packets with no VLAN tag (that is, untagged). All other packets with one or more VLAN tags (even those with priority tag only) are not processed and dropped. Users can use the service with svc- sap-type set to `null-star' to process and forward packets with one or more tags (including priority tag) on a null SAP.

Default null-star

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 7210 SAS on which this service is defined.

Values *service-id*: 1 — 2147483648

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

Values 1 — 2147483647

m-vpls — Specifies a management VPLS.

create — This keyword is mandatory while creating a VPLS service.

customer-vid vlan-id — Defines the dot1q VLAN ID to be specified while creating the local Dot1q SAP for svc-sap-type dot1q-preserve.

Values 1 — 4094

dot1q-preserve — Specifies that the allowed SAP in the service are Dot1q. The Dot1q ID is not stripped after packets matches the SAP.

Default null-star

null-star — Specifies that the allowed SAP in the service which can be Null SAP,dot1q Default SAP,Q.* SAP or Default QinQ SAP.

svc-sap-type- — Specifies the type of service and allowed SAPs in the service.

dot1q-range — Specifies that the access SAP in the service can use VLAN ranges as the SAP tags. The VLAN ranges are configured using the CLI command *configure> connection-profile*. On ingress of the access dot1q SAP using VLAN ranges, the received tag on the SAP is preserved. A VPLS service with svc-sap-type set to dot1q-range is used for Epipe emulation with G8032 for protection. For more information about the capabilities and restrictions, see Epipe Emulation using Dot1q VLAN range SAP in VPLS with G8032 on page 326.

r-vpls — Allows this VPLS instance to be associated with an IP interface to provide R-VPLS functionality. This parameter is supported on 7210 SAS-M, 7210 SAS-X and 7210 SAS-R6.

bgp

Syntax **bgp**

Context config>service>vpls

This command enables the context to configure the BGP related parameters to BGP AD.

block-on-mesh-failure

Syntax [**no**] **block-on-mesh-failure**

Context config>service>vpls>spoke-sdp
config>service>vpls>endpoint

Description This command enables blocking (brings the entity to an opererationally down state) after all configured SDPs or endpoints are in opererationally down state. This event is signalled to

corresponding T-LDP peer by withdrawing service label (status-bit-signaling non-capable peer) or by setting “PW not forwarding” status bit in T-LDP message (status-bit-signaling capable peer).

Default disabled

bpdu-translation

Syntax **bpdu-translation** {**auto** | **pvst** | **stp**}
no bpdu-translation

Context config>service>vpls>spoke-sdp
config>service>vpls>sap

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 **l2pt-termination** [cdp] [dtp] [pagp] [stp] [udld] [vtp]
no l2pt-termination

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command enables Layer 2 Protocol Tunneling (L2PT) termination on a given SAP. L2PT termination is supported for STP/CDP/DTP/PAGP/UDLD and VTP PDUs.

This feature can be enabled only if STP is disabled in the context of the given VPLS service.

Default no l2pt-termination

Parameters *cdp* — Specifies the Cisco discovery protocol.

dtp — Specifies the dynamic trunking protocol.

pagp — Specifies the port aggregation protocol.

stp — Specifies all spanning tree protocols: stp, rstp, mstp, pvst (default).

udld — Specifies unidirectional link detection.

vtp — Specifies the VLAN trunking protocol.

disable-aging

Syntax [**no**] **disable-aging**

Context config>service>vpls
config>service>vpls>spoke-sdp
config>service>vpls>sap
config>template>vpls-template
config>service>pw-template

Description This command disables MAC address aging across a VPLS service or on a VPLS service SAP.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the VPLS forwarding database (FDB). The **disable-aging** command turns off aging for local and remote learned MAC addresses.

When **no disable-aging** is specified for a VPLS, it is possible to disable aging for 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 [**no**] **disable-learning**

Context config>service>vpls
config>service>pw-template
config>template>vpls-template

Description This command disables learning of new MAC addresses in the VPLS forwarding database (FDB) for the service instance.

When **disable-learning** is enabled, new source MAC addresses will not be entered in the VPLS service forwarding database.

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.

endpoint

Syntax **endpoint** *endpoint-name* [**create**]
no endpoint

Context config>service>vpls

Description This command configures a service endpoint.

Parameters *endpoint-name* — Specifies an endpoint name up to 32 characters in length.

create — This keyword is mandatory while creating a service endpoint.

description

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

Context config>service>vpls>endpoint

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.

ignore-standby-signaling

Syntax [**no**] **ignore-standby-signaling**

Context config>service>vpls>endpoint
config>service>vpls>spoke-sdp

Description When this command is enabled, the node will ignore standby-bit received from TLDP peers for the given spoke SDP and performs internal tasks withot taking it into account.

This command is present at endpoint level as well as spoke SDP level. If the spoke SDP is part of the explicit-endpoint, it is not possible to change this setting at the spoke SDP level. The existing spoke SDP will become part of the explicit-endpoint only if the setting is not conflicting. The newly created spoke SDP which is a part of the given explicit-endpoint will inherit this setting from the endpoint configuration.

Default disabled

revert-time

Syntax **revert-time** *revert-time* **| infinite**
no revert-time

Context config>service>vpls>endpoint

Description This command configures the time to wait before reverting to primary spoke SDP.

In a regular endpoint the revert-time setting affects just the pseudowire defined as primary (precedence 0). For a failure of the primary pseudowire followed by restoration the revert-timer is started. After it expires the primary pseudowire takes the active role in the endpoint. This behavior does not apply for the case when both pseudowires are defined as secondary. For example, if the active secondary pseudowire fails and is restored it will stay in standby until a configuration change or a force command occurs.

Parameters *revert-time* — Specifies the time to wait, in seconds, before reverting back to the primary spoke SDP defined on this service endpoint, after having failed over to a backup spoke SDP.

Values 0 — 600

infinite — Specifying this keywork makes endpoint non-revertive.

split-horizon-group

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

Context config>service>pw-template

Description This command is used to create a new split horizon group for the VPLS instance. Traffic arriving on a SAP or spoke SDP within this split horizon group 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 instance. The same group-name can be re-used in different VPLS instances.

Note: In 7210-SAS devices, use of SAP or spoke-SDP Split-horizon group and Mesh-SDP are mutually exclusive.

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 SAP or Spoke-SDP belongs.

create — Mandatory keyword to create a split-horizon group.

static-mac

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

Context config>service>vpls>endpoint

Description This command assigns a static MAC address to the endpoint. In the FDB, the static MAC is then associated with the active spoke SDP.

Default none

Parameters *ieee-address* — Specifies the static MAC address to the endpoint.

Values 6-byte mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx). (Note: This value cannot be all zeros.)

create — This keyword is mandatory while creating a static MAC.

suppress-standby-signaling

Syntax [**no**] **suppress-standby-signaling**

Context config>service>vpls>endpoint

Description When this command is enabled, the pseudowire standby bit (with value 0x00000020) will not be sent to T-LDP peer when the given spoke is selected as a standby. This allows faster switchover as the traffic will be sent over this SDP and discarded at the blocking side of the connection. This is particularly applicable to multicast traffic.

Default enabled

propagate-mac-flush

Syntax [**no**] **propagate-mac-flush**

Context config>service>vpls

Description This command specifies whether MAC flush messages received from the given LDP are propagated to all spoke and mesh SDPs within the context of this VPLS service. The propagation will follow the split-horizon principle and any data-path blocking in order to avoid the looping of these messages.

Default no propagate-mac-flush

fdb-table-high-wmark

Syntax [**no**] **fdb-table-high-wmark** *high-water-mark*

Context config>service>vpls

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 0— 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 0— 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.

Parameters *table-size* — Specifies the maximum number of MAC entries in the FDB.

vsi-export

Syntax **vsi-export policy-name** *[policy-name...(up to 5 max)]*
no vsi-export

Context config>service>vpls>bgp-ad
config>service>vpls>bgp

Description This command specifies the name of the VSI export policies to be used for BGP auto-discovery, if this feature is configured in the VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

vsi-import

Syntax **vsi-import policy-name [policy-name...(up to 5 max)]**
no vsi-import

Context config>service>vpls>bgp-ad>vsi-id
config>service>vpls>bgp

Description This command specifies the name of the VSI import policies to be used for BGP auto-discovery, if this feature is configured in the VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.
The policy name list is handled by the SNMP agent as a single entity.

route-target

Syntax **route-target {ext-community|{[export ext-community][import ext-community]}}**
no route-target

Context config>service>vpls>bgp-ad
config>service>vpls>bgp

Description This command configures the route target (RT) component that will be signaled in the related MPBGP attribute to be used for BGP auto-discovery, if this feature is configured in the VPLS service.

If this command is not used, the RT is built automatically using the VPLS ID. The ext-comm can have the same two formats as the VPLS ID, a two-octet AS-specific extended community, IPv4 specific extended community.

The following rules apply:

- If BGP AD VPLS-id is configured & no RT is configured under BGP node - RT = VPLS-ID.
- If BGP AD VPLS-id is not configured then an RT value must be configured under BGP node. (this is the case when only BGP VPLS is configured)

- If BGP AD VPLS-id is configured and an RT value is also configured under BGP node, the configured RT value prevails

Parameters ***export*** *ext-community* — •Specify communities allowed to be sent to remote PE neighbors.

import *ext-community* — •Specify communities allowed to be accepted from remote PE neighbors.

pw-template-binding

Syntax **pw-template-binding policy-id** *[split-horizon-group group-name] [import-rt {extcommunity,...(up to 5 max)}]*
no pw-template-bind *policy-id*

Context config>service>vpls>bgp-ad
config>service>vpls>bgp

Description This command binds the advertisements received with the route target (RT) that matches the configured list (either the generic or the specified import) to a specific pw-template. If the RT list is not present the pw-template is used for all of them.

The pw-template-binding applies to BGP-AD, if this feature is configured in the VPLS service.

The tools perform commands can be used to control the application of changes in pw-template for BGP-AD.

The no form of the command removes the values from the configuration.

Default none

Parameters *policy-id* — Specifies an existing policy ID.

Values 1 — 2147483647

split-horizon-group group-name — The specified group-name overrides the split horizon group template settings.

import-rt ext-comm — Specify communities allowed to be accepted from remote PE neighbors. An extended BGP community in the type:x:y format. The value x can be an integer or IP address.

The type can be the target or origin. x and y are 16-bit integers.

Values target:{ip-addr:comm-val| 2byte-asnumber:ext-comm-val|4byte-asnumber:comm-val} ip-addr a.b.c.d

comm-val 0 — 65535
2byte-asnumber 0 — 65535
ext-comm-val 0 — 4294967295
4byte-asnumber 0 — 4294967295

route-distinguisher

Syntax **route-distinguisher [ip-addr:comm-val | as-number:ext-comm-val]**
no route-distinguisher

Context config>service>vpls>bgp-ad>vsi-id
config>service>vpls>bgp

Description This command configures the Route Distinguisher (RD) component that will be signaled in the MPBGP NLRI for L2VPN AFI. This value will be used for BGP-AD, if this feature is configured in the VPLS service.

If this command is not configured, the RD is automatically built using the BGP-AD VPLS ID. The following rules apply:

- If BGP AD VPLS-id is configured & no RD is configured under BGP node - RD = VPLS-ID.
- If BGP AD VPLS-id is not configured then an RD value must be configured under BGP node (this is the case when only BGP VPLS is configured).
- If BGP AD VPLS-id is configured and an RD value is also configured under BGP node, the configured RD value prevails Values and format (6 bytes, other 2 bytes of type is automatically generated)

Parameters *ip-addr:comm-val* — Specifies the IP address.

Values ip-addr a.b.c.d

comm-val 0 — 65535
as-number:ext-comm-val — Specifies the AS number and the
Values as-number 1 — 65535

ext-comm-val 0 — 4294967295

local-age

Syntax **local-age** *aging-timer*
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). MACs associated with a SAP are classified as local MACs, and MACs associated with are remote MACsQinQ / access uplink SAPs.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). The **local-age** timer specifies the aging time for local learned MAC addresses.

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

Default **local age 300** — Local MACs aged after 300 seconds.

Parameters *aging-timer* — 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** context. 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, as only those objects can be blocked by this feature.

The operation of this feature is the same on the SAP. For example, if a MAC address moves from SAP to SAP, one will be blocked to prevent thrashing.

mac-move will disable a VPLS port when the number of relearns detected has reached the number of relearns needed to reach the move-frequency in the 5-second interval. For example, when the move-frequency is configured to 1 (relearn per second) mac-move will disable one of the VPLS ports when 5 relearns were detected during the 5-second interval because then the average move-frequency of 1 relearn per second has been reached. This can already occur in the first second if the real relearn rate is 5 relearns per second or higher.

The **no** form of this command disables MAC move.

move-frequency

Syntax **move-frequency** *frequency*
no move-frequency

Context config>service>vpls>mac-move

Description This command 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 **no** form of the command reverts to the default value.

Default 2 (when mac-move is enabled). For example, 10 relearns in a 5 second period.

Parameters *frequency* — Specifies the rate, in 5-second intervals for the maximum number of relearns.

Values 1 — 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.

It is recommended that the retry-timeout value is larger or equal to 5s * cumulative factor of the highest priority port so that the sequential order of port blocking will not be disturbed by re-initializing lower priority ports.

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.

remote-age

Syntax **remote-age** *seconds*
no remote-age

Context config>service>vpls
config>template>vpls-template

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

Note: This command is applicable on 7210 SAS-M and 7210 SAS-T devices configured in network mode.

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. If the endpoint is configured within the VPLS and send-flush-on-failure is enabled, flush-all-from-me messages will be sent out only when all spoke SDPs associated with the endpoint go down.

This feature cannot be enabled on management VPLS.

Default no send-flush-on-failure

service-mtu

Note: This command is supported on 7210 SAS-M and 7210 SAS-T in Network mode.

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.

For i-VPLS and EPIPEs bound to a b-VPLS, the service-mtu must be at least 18 bytes smaller than the b-VPLS service MTU to accomodate the PBB header.

The **no** form of this command returns the default **service-mtu** for the indicated service type to the default value.

Note: To disable service MTU check execute the command no service-mtu-check. Disabling service MTU check allows the packets to pass to the egress if the packet length is lesser than or equal to the MTU configured on the port.

Default VPLS: 1514

The following table displays MTU values 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

The size of the MTU in octets, expressed as a decimal integer.

Values 1 — 9194

service-mtu-check

Note: This command is supported on 7210 SAS-M and 7210 SAS-T in Network mode.

Syntax [**no**] **service-mtu-check**

Context config>service>vpls

Description The **no** form of this command disables the service MTU checks. Disabling service MTU check allows the packets to pass to the egress if the packet length is lesser than or equal to the MTU configured on the port. The length of the packet sent from a SAP is limited only by the access port MTU. In case of a pseudowire the length of a packet is limited by the network port MTU (including the MPLS encapsulation).

Note: If TLDP is used for signaling,the configured value for service-mtu is used during a pseudowire setup.

Default enabled

split-horizon-group

Syntax [**no**] **split-horizon-group** [*group-name*] [**create**]

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.

Note: In 7210-SAS devices, use of SAP or Spoke-SDP Split-horizon group and Mesh-SDP are mutually exclusive. This command is supported on 7210 SAS-M and 7210 SAS-T in network mode. It is not available in 7210 SAS-M and 7210 SAS-T in access-uplink mode.

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 SAP or spoke-SDP belongs.

create — Mandatory keyword to create a split-horizon group.

root-guard

Syntax [**no**] **root-guard**

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>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

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.

vsi-id

Syntax **vsi-id**

Context config>service>vpls>bgp-ad

Description This command enables the context to configure the Virtual Switch Instance Identifier (VSI-ID).

prefix

Syntax **prefix low-order-vsi-id**
no prefix

Context config>service>vpls>bgp-ad>vsi-id

Description This command specifies the low-order 4 bytes used to compose the Virtual Switch Instance Identifier (VSI-ID) to use for NLRI in BGP auto-discovery in this VPLS service.

If no value is set, the system IP address will be used.

Default no prefix

Parameters *low-order-vsi-id* — Specifies a unique VSI ID.

Values 0— 4294967295

service-name

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

Context config>service>vpls

Description This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a given service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the 7750 SR, 7450 ESS and 7710 SR platforms.

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

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

allow-ip-int-binding

Syntax [no] **allow-ip-int-binding**

Context config>service>vpls

Description The allow-ip-int-binding command that sets a flag on the VPLS service that enables the ability to attach an IES IP interface to the VPLS service in order to make the VPLS service routable. When the allow-ip-int-binding command is not enabled, the VPLS service cannot be attached to an IP interface.

Please refer to the Virtual Private LAN Service on page 263 for VPLS Configuration Constraints for Enabling allow-ip-int-binding.
When attempting to set the allow-ip-int-binding VPLS flag, the system first checks to see if the correct configuration constraints exist for the VPLS service and the network ports. In Release 5.0 the following VPLS features must be disabled or not configured for the allow-ip-int-binding flag to set:

- SAP ingress QoS policies applied to the VPLS SAPs cannot have MAC match criteria defined
- The VPLS service type cannot be B-VPLS or M-VPLS and it cannot be an I-VPLS service bound to a B-VPLS context
- MVR from Routed VPLS and to another SAP is not supported

Once the VPLS allow-ip-int-binding flag is set on a VPLS service, the above features cannot be enabled on the VPLS service.

VPLS SERVICE NAME BOUND TO IP INTERFACE WITHOUT ALLOW-IP-INT-BINDING FLAG SET

In the event that a service name is applied to a VPLS service and that service name is also bound to an IP interface but the allow-ip-int-binding flag has not been set on the VPLS service context, the system attempt to resolve the service name between the VPLS service and the IP interface will fail. After the allow-ip-int-binding flag is successfully set on the VPLS service, either the service name on the VPLS service must be removed and reapplied or the IP interface must be re-initialized using the shutdown or no shutdown commands. This will cause the system to reattempt the name resolution process between the IP interface and the VPLS service.

The no form of the command resets the allow-ip-int-binding flag on the VPLS service. If the VPLS service currently has an IP interface from an IES service attached, the no allow-ip intbinding command will fail. Once the allow-ip-int-binding flag is reset on the VPLS service, the configuration restrictions associated with setting the flag are removed.

VPLS Interface Commands

Note: VPLS interface commands are supported only on 7210 SAS-M and 7210 SAS-T devices configured in access uplink mode.

interface

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

Context config>service>vpls

Description This command creates a logical IP routing interface for a VPLS service. Once created, attributes such as IP address and service access points (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 the VPLS service IDs. The IP interface created is associated with the VPLS management routing instance.This instance does not support routing.

Interface names are case-sensitive and must be unique within the group of defined IP interfaces defined for the network core router instance. Interface names in the dotted decimal notation of an IP address are not allowed. 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. Duplicate interface names can exist in different router instances.

Enter a new name to create a logical router interface. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, no default IP interface names are defined within the system. All VPLS IP interfaces must be explicitly defined in an enabled state.

The no form of this command removes the IP interface and the entire associated configuration. The interface must be administratively shutdown before issuing the no interface command.

For VPLS services, the IP interface must be shutdown before the SAP on that interface is removed.

For VPLS service, ping and traceroute are the only applications supported.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP.

An interface name:

- Should not be in the form of an IP address.
- Can be from 1 to 32 alphanumeric characters.
- If the string contains special characters (such as #,$,spaces), the entire string must be enclosed within double quotes.

If ip-int-name already exists within the service ID, the context changes to maintain that IP interface. If ip-int-name already exists within another service ID, an error occurs and the context does not change to that IP interface. If ip-int-name does not exist, the interface is created and the context is changed to that interface for further command processing.

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*}
address *ip-address mask*

Context config>service>vpls>interface

Description This command assigns an IP address and an IP subnet, to a VPLS IP router interface. Only one IP address can be associated with an IP interface. An IP address must be assigned to each VPLS IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the 7210 SAS.

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

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

Parameters *ip-address* — The IP address of the IP interface. The ip-address portion of the address command specifies the IP host address that will be used by the IP interface within the subnet.

This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

/ — The forward slash is a parameter delimiter and separates the ip-address portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the ipaddress, 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. The values allowed 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.

Values 1 — 16383

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.

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 a VPLS IP interface.

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

Default The system chassis MAC address.

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. 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 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.

VPLS STP Commands

stp

Syntax **stp**

Context config>service>vpls
config>service>vpls>sap
config>template>vpls-template

Description This command enables the context to configure the Spanning Tree Protocol (STP) parameters. 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.

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.

RSTP, however, can detect that the actual situation is different from what **edge-port** may indicate.

Initially, the value of the SAP or spoke SDP parameter is set to edge-port. This value 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 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
config>template>vpls-template>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 spends in the discarding and learning states when transitioning to the forwarding state.

The value of the forward-delay variable depends on the STP operating mode of the VPLS instance:

- in `rstp` or `mstp` mode, but only when the SAP has not fallen back to legacy STP operation, the value configured by the `hello-time` command is used;
- in all other situations, the value configured by the `forward-delay` command is used.

Default 15 seconds

Parameters *seconds* — The forward delay timer for the STP instance in seconds.

Values 4 — 30

hello-time

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

Context config>service>vpls>stp
config>template>vpls-template>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 428.

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
config>template>vpls-template>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. 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 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.

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. If a loop occurs, this parameter indicates the probability of a given port being assigned a forwarding state. (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.

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. If a loop occurs, this parameter indicates the probability of a given port being assigned a forwarding state.

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
config>template>vpls-template>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 | pmstp**}
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 283 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

pmstp — The PMSTP mode is only supported in VPLS services where the mVPLS flag is configured.

mst-instance

Syntax [**no**] **mst-instance** *mst-inst-number*

Context config>service>vpls>stp

Description This command creates the context to configure Multiple Spanning Tree Instance (MSTI) related parameters. MSTP supports “16” instances. The instance “0” is mandatory (by protocol) and cannot be created by the CLI. The software automatically maintains this instance.

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 values of the priority are only multiples of 4096 (4k). If a value is specified that is not a multiple of 4K, the value is replaced by the closest multiple of 4K(lower than the value entered).

The **no** form of this command sets the bridge-priority to its default value.

Default 32768 — All instances that are created by the **vlan-range** command do not have explicit definition of bridge-priority and will inherit the default value.

Parameters *bridge-priority* — Specifies the priority of this specific 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— 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 if their configured MST-region name, MST-revision, and VLAN-to-instance are 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 path cost is used by STP to calculate the path cost to the root bridge. The path cost in BPDUs received on the root port is incremented with the configured path cost for that SAP. When BPDUs are sent out other egress SAPs or spoke SDPs, the newly calculated root path cost is used.

STP suggests that the path cost is defined as a function of the link bandwidth. Since SAPs are controlled by complex queuing dynamics, in the 7210 SAS the STP path cost is a purely static configuration.

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

path-cost — The path cost for the SAP or spoke SDP.

Values 1 — 200000000 (1 is the lowest cost)

Default 10

port-num

Syntax [**no**] **port-num** *virtual-port-number*

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

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
config>template>vpls-template>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 within the STP instance.

STP computes the actual priority by taking the input value and masking out the lower four bits.The result is the value that is stored in the priority parameter. For instance, if a value of 0 is entered, masking out the lower 4 bits results in a parameter value of 0. If a value of 255 is entered, the result is 240.

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

Default 128

Parameters *stp-priority* — The STP priority value for the SAP . Allowed values are integer in the range of 0 to 255, 0 being the highest priority. The actual value used for STP priority (and stored in the configuration) 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**] [**eth-ring** *ring-index*] (for 7210 SAS-M and 7210 SAS-T in Network mode) **sap** sap-id [g8032-shg-enable] [eth-ring *ring-index*] [**create**] (for 7210 SAS-M and 7210 SAS-T in Access uplink mode)
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 7210 SAS. Each SAP must be unique.

A physical port can have only one SAP to be part of one service. Multiple SAPS can be defined over a physical port but each of these SAPs should belong to a different service.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP 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.

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.

This command is also used to create a Ring APS Control SAP or a Data SAP whose traffic is protected by a Ring APS Instance.

No SAPs are defined.

Special Cases A default SAP has the following format: port-id:*. This type of SAP is supported only on Ethernet MDAs and its creation is allowed only in the scope of Layer 2 services (Epipe and VPLS).The 7210 SAS supports explicit null encapsulation for VPLS service.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 987 for command syntax.

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

g8032-shg-enable — Platforms Supported - 7210 SAS-M (access-uplink mode) and 7210 SAS-T. This command must only be used with the SAPs created in the service for the virtual channel on the interconnection nodes in a topology that uses multiple rings. This command creates a split-

horizon group to ensure that Sub-Ring control messages from the major ring are only passed to the Sub-Ring control service.

eth-ring — The keyword to create an instance of a Ring APS Control SAP or a Data SAP whose traffic is protected by a Ring APS Instance.

ring-index — Specifies the ring index of the Ethernet ring.

split-horizon-group *group-name* **—** Specifies the name of the split horizon group to which the SAP belongs.

discard-unknown-source

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

Context config>service>vpls>sap

Description When this command is enabled, packets received on a SAP 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 446) 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-source**

config>service>vpls

ETH-CFM Service Commands

eth-cfm

Syntax **eth-cfm**

Context config>service>vpls
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>vpls>sap

Description This command enables the context to configure ETH-CFM parameters.

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {up | **down**}] **primary-vlan-enable** [**vlan** vlan-id]
no mep *mep-id* **domain** *md-index* **association** *ma-index*

Context config>service>vpls>mesh-sdp>eth-cfm
config>service>vpls>sap>eth-cfm

Description This command configures the ETH-CFM maintenance endpoint (MEP).

mep-id — Specifies the maintenance association end point identifier.

Values 1 — 8191

md-index — Specifies the maintenance domain (MD) index value.

Values 1 — 4294967295

ma-index — Specifies the MA index value.

Values 1 — 4294967295

direction up| down — Indicates the direction in which the maintenance association (MEP) faces on the bridge port. Direction is not supported when a MEP is created directly under the vpls>eth-cfm construct (vMEP).

down — Sends ETH-CFM messages away from the MAC relay entity.

up — Sends ETH-CFM messages towards the MAC relay entity.

primary-vlan-enable — Provides a method for linking the MEP with the primary VLAN configured under the bridge-identifier for the MA. MEPs can not be changed from or to primary vlan functions. This must be configured as part of the creation step and can only be changed by deleting the MEP and recreating it. Primary VLANs are only supported under Ethernet SAPs.

ais-enable

Syntax [**no**] **ais-enable**

Context config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command enables the generation and the reception of AIS messages.

client-meg-level

Syntax **client-meg-level** [[*level* [*level* ...]]
no client-meg-level

Context config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable
config>service>epipe>spoke-sdp>eth-cfm>mep>ais-enable

Description This command configures the client maintenance entity group (MEG) level(s) to use for AIS message generation. Up to 7 levels can be provisioned with the restriction that the client MEG level must be higher than the local MEG level.

Parameters *level* — Specifies the client MEG level.

Values 1 — 7

Default 1

interval

Syntax **interval** {**1** | **60**}
no interval

Context config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable
config>service>epipe>spoke-sdp>eth-cfm>mep>ais-enable

Description This command specifies the transmission interval of AIS messages in seconds.

Parameters **1** | **60 —** The transmission interval of AIS messages in seconds.

Default 1

priority

Syntax **priority** *priority-value*
no priority

Context config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable
config>service>epipe>spoke-sdp>eth-cfm>mep>ais-enable

Description This command specifies the priority of AIS messages originated by the node.

Parameters *priority-value* — Specify the priority value of the AIS messages originated by the node.

ccm-enable

Syntax [**no**] **ccm-enable**

Context config>service>vpls>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>mesh-sdp>mep
config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command enables the generation of CCM messages.

The **no** form of the command disables the generation of CCM messages.

ccm-ltm-priority

Syntax **ccm-ltm-priority** *priority*
no ccm-ltm-priority

Context config>service>vpls>sap>eth-cfm>mep
config>service>vpls>mesh-sdp>mep
config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command specifies the priority value for CCMs and LTMs transmitted by the MEP.

The **no** form of the command removes the priority value from the configuration.

Default The highest priority on the bridge-port.

Parameters *priority* — Specifies the priority of CCM and LTM messages.

Values 0 — 7

eth-test-enable

Syntax [**no**] **eth-test-enable**

Context config>service>vpls>spoke-sdp>eth-cfm>mep

Description For ETH-test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands:

oam eth-cfm eth-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*] [data-length *data-length*]

A check is done for both the provisioning and test to ensure the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the CLI and SNMP will indicate the problem.

test-pattern

Syntax **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
no test-pattern

Context config>service>vpls>sap>eth-cfm>mep>eth-test-enable
config>service>vpls>mesh-sdp>eth-cfm>mep>eth-test-enable

Description This command configures the test pattern for eth-test frames.

The **no** form of the command removes the values from the configuration.

Parameters **all-zeros —** Specifies to use all zeros in the test pattern.

all-ones — Specifies to use all ones in the test pattern.

crc-enable — Generates a CRC checksum.

Default all-zeros

fault-propagation-enable

Syntax **fault-propagation-enable {use-if-tlv | suspend-ccm}**
no fault-propagation-enable

Context config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command configures the fault propagation for the MEP.

Parameters *use-if-tlv* — Specifies to use the interface TLV.

suspend-ccm — Specifies to suspend the continuity check messages.

low-priority-defect

Syntax **low-priority-defect** {**allDef**|**macRemErrXcon**|**remErrXcon**|**errXcon**|**xcon**|**noXcon**}

Context config>service>vpls>mesh-sdp>eth-cfm>mep

config>service>epipe>sap>eth-cfm>mep>eth-test-enable

config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm.

Default macRemErrXcon

Values

allDef	DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
macRemErrXcon	Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
remErrXcon	Only DefRemoteCCM, DefErrorCCM, and DefXconCCM

errXcon	Only DefErrorCCM and DefXconCCM
xcon	Only DefXconCCM; or
noXcon	No defects DefXcon or lower are to be reported

mac-address

Syntax **mac-address** *mac-address*
no mac-address

Context config>service>vpls>mesh-sdp>eth-cfm>mep

config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command specifies the MAC address of the MEP.

The **no** form of this command reverts the MAC address of the MEP back to that of the port (if the MEP is on a SAP) or the bridge (if the MEP is on a spoke).

Parameters *mac-address* — Specifies the MAC address of the MEP.

Values 6-byte mac-address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MEP. Must be unicast. Using the all zeros address is equivalent to the no form of this command.

one-way-delay-threshold

Syntax **one-way-delay-threshold** *seconds*

Context config>service>vpls>sap>eth-cfm>mep

config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command enables/disables eth-test functionality on MEP.

Parameters *seconds* — Specifies the one way delay threshold, in seconds.

Values 0..600

Default 3

tunnel-fault

Syntax **tunnel-fault** {**accept | ignore**}

Context config>service>vpls>eth-cfm
config>service>vpls>sap>eth-cfm

Description Allows the individual service SAPs to react to changes in the tunnel MEP state. When tunnel-fault accept is configured at the service level, the SAP will react according to the service type, EPIPE will set the operational flag and VPLS, IES and VPRN SAP operational state will become down on failure or up on clear. This command triggers the OAM mapping functions to mate SAPs and bindings in an EPIPE service as well as setting the operational flag. If AIS generation is the requirement for the

EPIPE services this command is not required. See the command ais-enable under epipe>sap>eth-cfm>ais-enable for more details. This works in conjunction with the tunnel-fault accept on the individual SAPs. Both must be set to accept to react to the tunnel MEP state. By default the service level command is “ignore” and the sap level command is “accept”. This means simply changing the service level command to “accept” will enable the feature for all SAPs. This is not required for EPIPE services that only wish to generate AIS on failure.

Parameters **accept —** Share fate with the facility tunnel MEP

ignore — Do not share fate with the facility tunnel MEP

Default **ignore** (Service Level)

accept (SAP Level for EPIPE and VPLS)

limit-mac-move

Syntax **limit-mac-move** [**blockable** | **non-blockable**]
no limit-mac-move

Context config>service>vpls>spoke-sdp
config>service>vpls>sap

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 **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

config>service>pw-template

Description This command disables re-learning of MAC addresses on other mesh SDPs within the VPLS. The MAC address remains attached to a given Mesh for duration of its age-timer.
The age of the MAC address entry in the FIB is set by the age timer. If mac-aging is disabled on a given VPLS service, any MAC address learned on a mesh with mac-pinning enabled remains in the FIB on this mesh forever. Every event that otherwise results in re-learning is logged (MAC address; original - mesh SDP; new - mesh SDP).

Default MAC pinning is not enabled by default.

max-nbr-mac-addr

Syntax **max-nbr-mac-addr** *table-size*
no max-nbr-mac-addr

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>endpoint

config>service>pw-template

Description This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this SAP, spoke SDP or endpoint.

When the configured limit has been reached, and discard-unknown-source has been enabled for this SAP or spoke SDP (see discard-unknown-source on page 439), 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 — 30719

statistics

Syntax **statistics**

Context config>service>vpls>sap

Description This command enables the context to configure the counters associated with SAP ingress and egress.

ingress

Syntax **ingress**

Context config>service>epipe>sap>statistics
config>service>vpls>sap>statistics

Description This command enables the context to configure the ingress SAP statistics counter.

counter-mode

Syntax **counter-mode {in-out-profile-count| forward-drop-count}**

Context config>service>epipe>sap>statistics>ingress
config>service>vpls>sap>statistics>ingress

Description This command allows the user to set the counter mode for the counters associated with sap ingress meters (a.ka. policers). A pair of counters is available with each meter. These counters count different events based on the counter mode value.

Note: The counter mode can be changed if an accounting policy is associated with a SAP. If the counter mode is changed the counters associated with the meter are reset and the counts are cleared. If an accounting policy is in use when the counter-mode is changed a new record will be written into the current accounting file.

Execute the following sequence of commands to ensure a new accounting file is generated when the counter-mode is changed:

1. Execute the command **config>service>epipe/vpls>sap> no collect-stats**, to disable writing of accounting records.
2. Change the counter-mode to the desired value, execute the command **config>service>epipe/vpls>sap>counter-mode {in-out-profile-count| forward-drop-count}**.
3. Execute the command **config>service>epipe/vpls>sap> collect-stats**, to enable writing of accounting records.

The **no** form of the command restores the counter mode to the default value.

Default in-out-profile-count

Parameters **forward-drop-count —** If the counter mode is specified as "forward-drop-count", one counter counts the forwarded packets and octets received on ingress of a SAP and another counts the dropped packets. The forwarded count is the sum of in-profile and out-of-profile packets/octets received on SAP ingress. The dropped count is count of packets/octets dropped by the policer. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. The in-profile count and out-of-profile count is not individually available when operating in this mode.

in-out-profile-count — If the counter mode is specified as "in-out-profile-count", one counter counts the total in-profile packets and octets received on ingress of a SAP and another counts the total out-of-profile packets and octets received on ingress of a SAP. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. Dropped counts are not maintained in hardware when this mode is used. It is obtained by subtracting the sum of in-profile count and out-of-profile count from the total SAP ingress received count and displayed.

static-mac

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

Context config>service>vpls>sap
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 definitions on one edge device are not propagated to other edge devices participating in the VPLS instance, that is, each edge device has an independent forwarding database for the VPLS.

Only one static MAC entry (local or remote) can be defined per MAC address per VPLS instance.

By default, no static MAC address entries are defined for the SAP.

The **no** form of this command deletes the static MAC entry with the specified MAC address associated with the SAP from the VPLS forwarding database.

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

create — This keyword is mandatory when specifying a static MAC address.

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 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.

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 375.

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 Filter and QoS Policy Commands

egress

Syntax **egress**

Context config>service>vpls>sap

Description This command enables the context to configure egress filter policies.

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.

filter

Syntax **filter ip** *ip-filter-id*
filter ipv6 *ipv6-filter-id*
filter mac *mac-filter-id*

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress
config>service>vpls>spoke-sdp>egress

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.

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

qos

Syntax **qos** *policy-id*
no qos

Context config>service>vpls>sap>ingress

Description This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP) or IP interface.

QoS ingress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the *policy-id* does not exist, an error will be returned.

The **qos** command is used to associate ingress apolicies. The **qos** command only allows ingress policies to be associated on SAP ingress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a given type will return an error.

By default, if no specific QoS policy is associated with the SAP for ingress , so the default QoS policy is used.

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

policy-id — The ingress policy ID to associate with SAP on ingress. The policy ID must already exist.

Values 1 — 65535

aggregate-meter-rate

Syntax **aggregate-meter-rate** *rate-in-kbps* [**burst** *burst-in-kbits*]
no aggregate-meter-rate

Context config>service> vpls> sap> ingress

config>service>epipe> sap> ingress

Description This command allows the user to configure the SAP aggregate policer. The rate of the SAP aggregate policer must be specified by the user. The user can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the traffic on different FCs and determines the destination of the packet. The packet is either forwarded to an identified profile or dropped.

Note: The sum of CIR of the individual FCs configured under the SAP cannot exceed the PIR rate configured for the SAP. Though the 7210 SAS software does not block this configuration, it is not recommended for use.

The table below provides information about the final disposition of the packet based on the operating rate of the per FC policer and the per SAP aggregate policer:

Per FC meter Operating Rate	Per FC Assigned Color	SAP aggregate meter Operating Rate	SAP aggregate meter color	Final Packet Color
Within CIR	Green	Within PIR	Green	Green or In-profile
Within CIR*	Green	Above PIR	Red	Green or In-profile
Above CIR, Within PIR	Yellow	Within PIR	Green	Yellow or Out-of-Profile
Above CIR, Within PIR	Yellow	Above PIR	Red	Red or Dropped
Above PIR	Red	Within PIR	Green	Red or Dropped
Above PIR	Red	Above PIR	Red	Red or Dropped

Table 23: Final Disposition of the packet based on per FC and per SAP policer or meter.

Note*: The row number 2 in the above table is not recommended for use. For more information on this, see the Note in the "**aggregate-meter-rate**" description.

When the SAP aggregate policer is configured, per FC policer can be only configured in "trtcm2" mode (RFC 4115).

Note: The meter modes "srtcm" and "trtcm1" are used in the absence of an aggregate meter.

The SAP ingress meter counters increment the packet or octet counts based on the final disposition of the packet.

If ingress Frame-based accounting is used, the SAP aggregate meter rate accounts for the Ethernet frame overhead. The system accounts for 12 bytes of IFG and 8 bytes of start delimiter.

The **no** form of the command removes the aggregate policer from use.

Default no aggregate-meter-rate

Parameters *rate-in-kbps* — Specifies the rate in kilobits per second.

Values 0 — 20000000 | max

Default max

burst <*burst-in-kilobits*> — Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 —2146959

Default 512

accounting-policy

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

Context config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>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 — 99

collect-stats

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

Context config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>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 cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent

collect-stats command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default no collect-stats

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 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 router. If two *sdp-id* bindings terminate on the same router, an error occurs and the second SDP is binding is rejected.

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

vlan — 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 mesh SDP bindings.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [create] [**split-horizon-group** *group-name*]
no spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [create]

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 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

create — This keyword is mandatory while creating a spoke SDP.

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. (hex 5)

split-horizon-group *group-name* **—** Specifies the name of the split horizon group to which the SDP belongs.

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

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.

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.

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

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.

fast-leave

Syntax [**no**] **fast-leave**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>pw-template>igmp-snooping

Description This command enables fast leave. When IGMP fast leave processing is enabled, the 7210 SAS M will immediately remove a SAP or SDP from the multicast group when it detects an IGMP “leave” on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a 'leave'

from the forwarding table without first sending out group-specific queries to the SAP or SDP, and 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** *service-id*
no from-vpls

Context config>service>vpls>sap>igmp-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 *service-id* — Specifies the MVR VPLS from which multicast channels should be copied into this SAP.

Values *service-id*: 1 — 2147483648

group

Syntax [**no**] **group** *grp-address*

Context config>service>vpls>sap>igmp-snooping>static
config>service>vpls>spoke-sdp>snooping>static
config>service>vpls>mesh-sdp>snooping>static

This command adds a static multicast group as a (*, g). When a static IGMP group is added, multicast data for that (*,g) is forwarded to the specific SAP 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>igmp-snooping>mvr

Description This command identifies filter policy of multicast groups to be applied to this VPLS entity. The

sources of the multicast traffic must be a member of the VPLS.

The **no** form of the command removes the policy association from the VPLS configuration.

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.

force-vlan-vc-forwarding

Syntax [**no**] **force-vlan-vc-forwarding**

Context config>service>epipe>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>pw-template

This command forces vc-vlan-type forwarding in the data path for spoke/mesh SDPs which have either vc-type. This comand is not allowed on vlan-vc-type SDPs.

The **no** form of this command sets default behavior.

Default disabled

igmp-snooping

Syntax **igmp-snooping**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>pw-template

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>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping
config> service>vpls> mesh-sdp>igmp-snooping
config>service>pw-template>igmp-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 or SDP 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. Values can be string up to 32 characters long of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. These policies are configured in the **config**>**router**> **policy-options** context The router policy must be defined before it can be imported.

last-member-query-interval

Syntax **last-member-query-interval** *tenths-of-seconds*
no last-member-query-interval

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>pw-template>igmp-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

max-num-groups

Syntax **max-num-groups** *count*
no max-num-groups

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>pw-template>igmp-snooping

Description This command defines the maximum number of multicast groups that can be joined on this SAP or SDP. If the node receives an IGMP join message that would exceed the configured number of groups, the request is ignored.

Default no max-num-groups

Parameters *count* — Specifies the maximum number of groups that can be joined on this SAP or SDP.

mrouter-port

Syntax [**no**] **mrouter-port**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping

Description This command specifies whether a multicast router is attached behind this SAP.

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 (v1or v2) can only be determined after an initial query has been received. Until such time no reports are sent on the SAP or SDP, even if mrouter-port is enabled.

If the **send-queries** command is enabled on this SAP, the mrouter-port parameter can not be set.

Default no mrouter-port

mvr

Syntax mvr

Context config>service>vpls>igmp-snooping
config>service>vpls>sap>igmp-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>igmp-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping

config>service>vpls>mesh-sdp>igmp-snooping
config>service>pw-template>igmp-snooping

Description This command configures the IGMP query interval. If the **send-queries** command is enabled, this parameter specifies the interval between two consecutive general queries sent by the system on this SAP or SDP. The configured query-interval must be greater than the configured query-response-interval. If send-queries is not enabled on this SAP or SDP, the configured query-interval value is ignored.

Default 125

Parameters *seconds* — The time interval, in seconds, that the router transmits general host-query messages.

Values 2 — 1024

query-src-ip

Note: This command is supported only on 7210 SAS-M and 7210 SAS-T devices configured in Network mode.

Syntax **query-src-ip** *ip-address*
no query-src-ip

Context config>service>vpls>igmp-snooping

Description This command configures the IP source address used in IGMP queries.

query-response-interval

Syntax **query-response-interval** *seconds*

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>pw-template>igmp-snooping

Description This command configures the IGMP query response interval. If the **send-queries** command is enabled, this parameter specifies the maximum response time advertised in IGMP queries.

The configured query-response-interval must be smaller than the configured query-interval.

If send-queries is not enabled on this SAP or SDP, the configured query-response-interval value is ignored.

Default 10

Parameters *seconds* — Specifies the length of time to wait to receive a response to the host-query message from the host.

Values 1 — 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>igmp-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>pw-template>igmp-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 **config>service>vpls>sap>igmp-snooping**: 2— 7
config>service>vpls>igmp-snooping: 1 — 255
config>service>vpls>spoke->sdp>igmp-snooping: 2— 7
config>service>vpls>mesh-sdp>igmp-snooping: 2— 7

precedence

Syntax **precedence** *precedence-value***| primary**
no precedence

Context config>service>vpls>spoke-sdp

Description This command configures the spoke SDP precendence.

Default 4

Parameters *precedence-value* — Specify the spoke SDP precedence.

Values 0 — 4

primary — Specifies that the prececence is primary.

propagate-mac-flush

Syntax [**no**] **propagate-mac-flush**

Context config>service>vpls

Description This command specifies whether MAC flush messages received from the given LDP are propagated to all spoke and mesh SDPs within the context of this VPLS service. The propagation will follow the split-horizon principle and any data-path blocking in order to avoid the looping of these messages.

Default no propagate-mac-flush

send-queries

Syntax [**no**] **send-queries**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>pw-template>igmp-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 will be the version of the querier.

Default no send-queries

starg

Syntax [**no**] **starg**

Context config>service>vpls>sap>igmp-snooping>static>group
config>service>vpls>spoke-sdp>igmp-snooping>static>group
config>service>vpls>mesh-sdp>igmp-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.

The **no** form of the command removes the starg entry from the configuration.

Default no starg

static

Syntax **static**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-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) 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>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>snooping>static
config>service>pw-template>igmp-snooping

Description This command specifies the version of IGMP which is running on this SAP or SDP. This object can be used to configure a router capable of running either value. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN.

When the **send-query** command is configured, all 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 effectThe version used on that SAP orwill be the version of the querier.

Note: IGMP V3 is supported only on 7210 SAS-M devices configured in access-uplink mode. IGMP V3 is not supported on 7210 SAS-M and 7210 SAS-T devices configured in network mode.

Parameters *version* — Specify the IGMP version.

to-sap

Syntax **to-sap** *sap-id*
no to-sap

Context config>service>vpls>sap>igmp-snooping>mvr

Description This command configures the SAP to which the multicast data needs to be copied.

In some scenarios, the multicast traffic should not be copied from the MVR VPLS to the SAP on which the IGMP message was received (standard MVR behaviour) but to another SAP.

Default no to-sap

Parameters *sap-id* — Specifies the SAP to which multicast channels should be copied.

IEEE 802.1ah Provider Backbone Bridging

In This Chapter

Note: PBB is supported on 7210 SAS-M and 7210 SAS-T devices configured in Network mode devices.

This chapter provides information about Provider Backbone Bridging (PBB), process overview, and implementation notes.

Topics in this chapter include:

IEEE 802.1ah Provider Backbone Bridging (PBB) Overview

IEEE 802.1ah draft standard (IEEE802.1ah), also known as Provider Backbone Bridges (PBB), defines an architecture and bridge protocols for interconnection of multiple Provider Bridge Networks (PBNs - IEEE802.1ad QinQ networks). PBB is defined in IEEE as a connectionless technology based on multipoint VLAN tunnels. IEEE 802.1ah employs Provider MSTP as the core control plane for loop avoidance and load balancing. As a result, the coverage of the solution is limited by STP scale in the core of large service provider networks. The 7210 SAS-M and 7210 SAS-T in network mode supports a native PBB Ethernet backbone deployment.

The IEEE model for PBB is organized around a B-component handling the provider backbone layer and an I-component concerned with the mapping of Customer or Provider Bridge (QinQ) domain (for example, MACs, VLANs) to the provider backbone (for example, B-MACs, B-VLANs), that is, the I-component contains the boundary between the Customer and Backbone MAC domains. PBB encapsulates customer payload in a provider backbone Ethernet header, providing for Customer MAC hiding capabilities. With PBB, 7210 devices can be used for tier-1/2 aggregation, encapsulating customer service frames in PBB, allowing the PE-rs devices deployed in the metro core to be aware of only provider MAC addresses and for metro service scaling.

7210 devices fully support only native PBB deployment. They do not support the integrated PBB VPLS model. In particular, 7210 devices do not support use of SDPs in PBB services.

PBB Features

Integrated PBB-VPLS Solution

HVPLS introduced a service-aware device in a central core location in order to provide efficient replication and controlled interaction at domain boundaries. The core network facing provider edge (N-PE) devices have knowledge of all VPLS services and customer MAC addresses for local and related remote regions resulting in potential scalability issues as depicted in Figure 58.

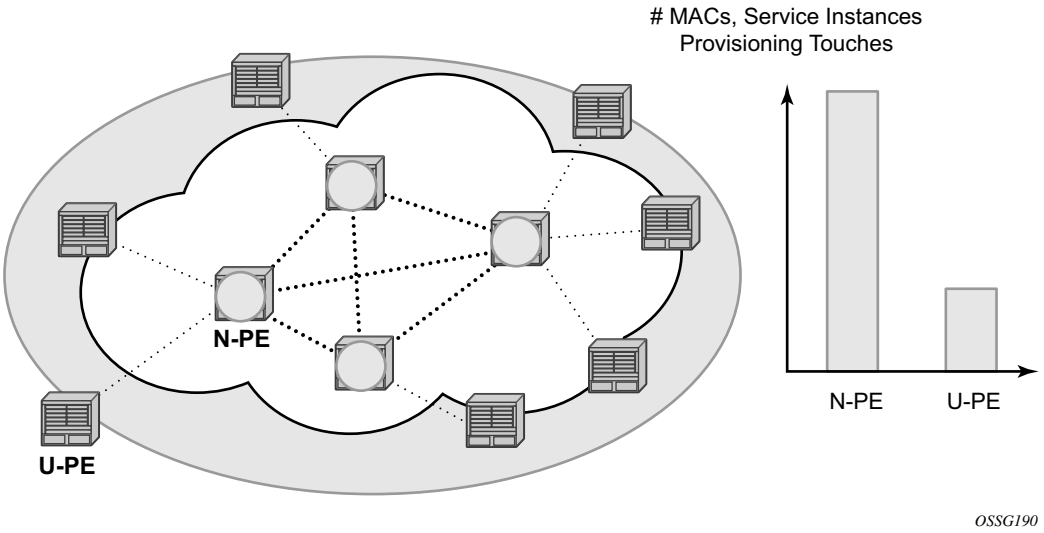

Figure 58: Large HVPLS Deployment

In a large VPLS deployment, it is important to improve the stability of the overall solution and to speed up service delivery. These goals are achieved by reducing the load on the N-PEs and respectively minimizing the number of provisioning touches on the N-PEs.

The integrated PBB-VPLS model introduces an additional PBB hierarchy in the VPLS network to address these goals as depicted in Figure 59.

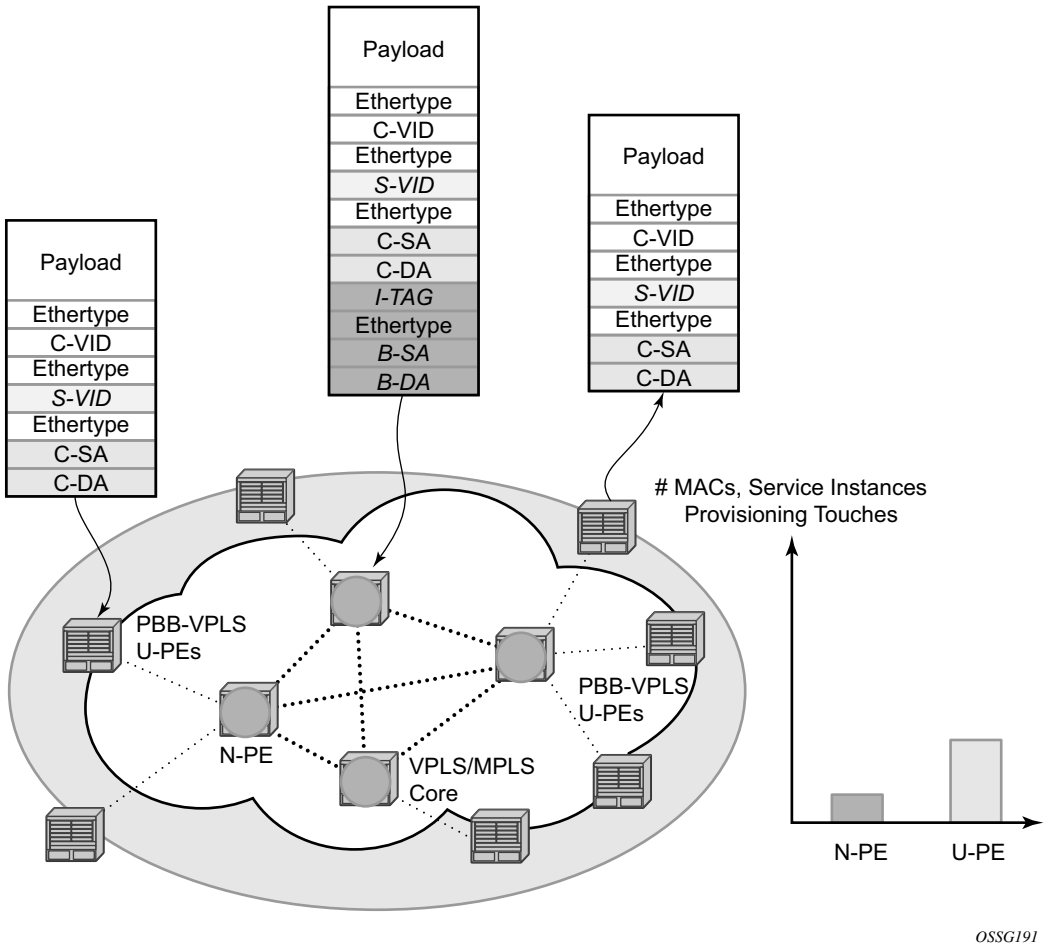

Figure 59: Large PBB-VPLS Deployment

PBB encapsulation is added at the user facing PE (U-PE) to hide the customer MAC addressing and topology from the N-PE devices. The core N-PEs need to only handle backbone MAC addressing and do not need to have visibility of each customer VPN. As a result, the integrated PBB-VPLS solution decreases the load in the N-PEs and improves the overall stability of the backbone.

In Figure 59, 7210 devices can only be used as U-PEs supporting only native Ethernet PBB services.

PBB Technology

IEEE 802.1ah specification encapsulates the customer or QinQ payload in a provider header as shown in Figure 60.

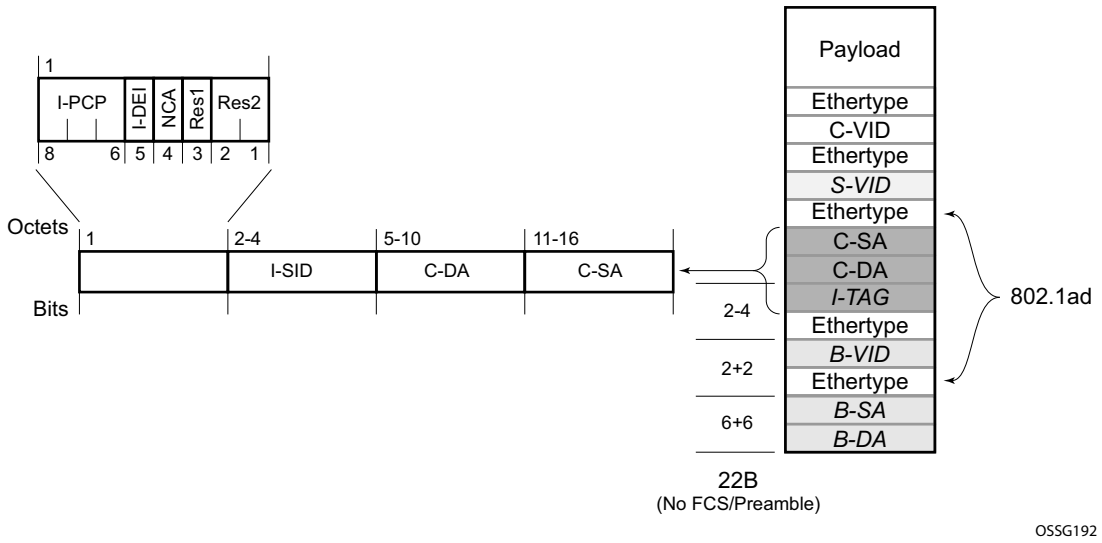

Figure 60: QinQ Payload in Provider Header Example

PBB adds a regular Ethernet header where the B-DA and B-SA are the backbone destination and respectively, source MACs of the edge U-PEs. The backbone MACs (B-MACs) are used by the core N-PE devices to switch the frame through the backbone.

A special group MAC is used for the backbone destination MAC (B-DA) when handling an unknown unicast, multicast or broadcast frame. This backbone group MAC is derived from the I-service instance identifier (ISID) using the rule: a standard group OUI (01-1E-83) followed by the 24 bit ISID coded in the last three bytes of the MAC address.

The BVID (backbone VLAN ID) field is a regular DOT1Q tag and controls the size of the backbone broadcast domain.

The following ITAG (standard Ether-type value of 0x88E7) has the role of identifying the customer VPN to which the frame is addressed through the 24 bit ISID.

PBB Mapping to Existing VPLS Configurations

The IEEE model for PBB is organized around a B-component handling the provider backbone layer and an I-component concerned with the mapping of the customer/provider bridge (QinQ) domain (MACs, VLANs) to the provider backbone (B-MACs, B-VLANs). For example, the I-component contains the boundary between the customer and backbone MAC domains.

Alcatel-Lucent's implementation is extending the IEEE model for PBB to allow support for MPLS pseudowires using a chain of two VPLS context linked together as depicted in Figure 61.

7210 does not support MPLS pseudowires in a PBB B-component and PBB I-component.

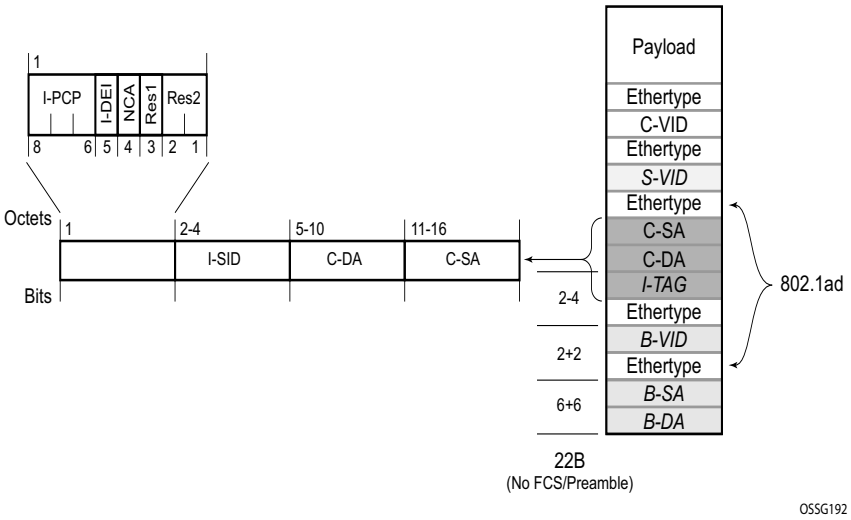

Figure 61: PBB Mapping to VPLS Constructs

Note: I-PW and B-PW are not supported on 7210 SAS devices.

A VPLS context is used to provide the backbone switching component. The white circle marked B, referred to as backbone-VPLS (B-VPLS) operates on backbone MAC addresses providing a core multipoint infrastructure that may be used for one or multiple customer VPNs. Alcatel-Lucent's B-VPLS implementation allows the use of native PBB infrastructures.

Note: 7210 implementation allows the use of only native PBB over Ethernet infrastructures.

Another VPLS context (I-VPLS) can be used to provide the multipoint I-component functionality emulating the ELAN service (refer to the triangle marked "I" in Figure 61). Similar to B-VPLS, I-VPLS inherits from the regular VPLS and native Ethernet (SAPs) handoffs accommodating this way different types of access: for example, direct customer link, QinQ or HVPLS.

In order to support PBB ELINE (point-to-point service), the use of an Epipe as I-component is allowed. All Ethernet SAPs supported by a regular Epipe are also supported in the PBB Epipe.

Note: 7210 implementation allows the use of only native PBB over Ethernet infrastructures.

SAP Support

PBB B-VPLS

- SAPs
  - → Ethernet DOT1Q is supported — This is applicable to most PBB use cases, for example, one backbone VLAN ID used for native Ethernet tunneling.
  - → Ethernet null is supported — This is supported for a direct connection between PBB PEs, for example, no BVID is required.
  - → Default SAP types are blocked in the CLI for the B-VPLS SAP.
- The following rules apply to the SAP processing of PBB frames:
  - → For “transit frames” (not destined to a local BMAC), there is no need to process the ITAG component of the PBB Frames. Regular Ethernet SAP processing is applied to the backbone header (BMACs and BVID).
  - → If a local I-VPLS instance is associated with the B-VPLS, “local frames” originated/terminated on local I-VPLS(s) are PBB encapsulated/de-encapsulated using the **pbb-etype** = 0x88e7.

PBB I-VPLS

- Port Level
  - → All existing Ethernet encapsulation types are supported (for example, null, dot1q, qinq).
- SAPs
  - → The I-VPLS SAPs can co-exist on the same port with SAPs for other business services, for example, VLL, VPLS SAPs.
  - → All existing Ethernet encapsulation are supported: null, dot1q, qinq.

Existing SAP processing rules still apply for the I-VPLS case; the SAP encapsulation definition on Ethernet ingress ports defines which VLAN tags are used to determine the service that the packet belongs to:

- Null encap defined on ingress — Any VLAN tags are ignored and the packet goes to a default service for the SAP.
- Dot1q encap defined on ingress — only first VLAN tag is considered;
- Qinq encap defined on ingress — both VLAN tags are considered; wildcard support for the inner VLAN tag

- For dot1q/qinq encapsulations, traffic encapsulated with VLAN tags for which there is no definition is discarded.
- Note that any VLAN tag used for service selection on the I-SAP is stripped before the PBB encapsulation is added. Appropriate VLAN tags are added at the remote PBB PE when sending the packet out on the egress SAP.

PBB Packet Walkthrough

This section describes the walkthrough for a packet that traverses the B-VPLS and I-VPLS instances using the example of a unicast frame between two customer stations as depicted in the following network diagram Figure 62.

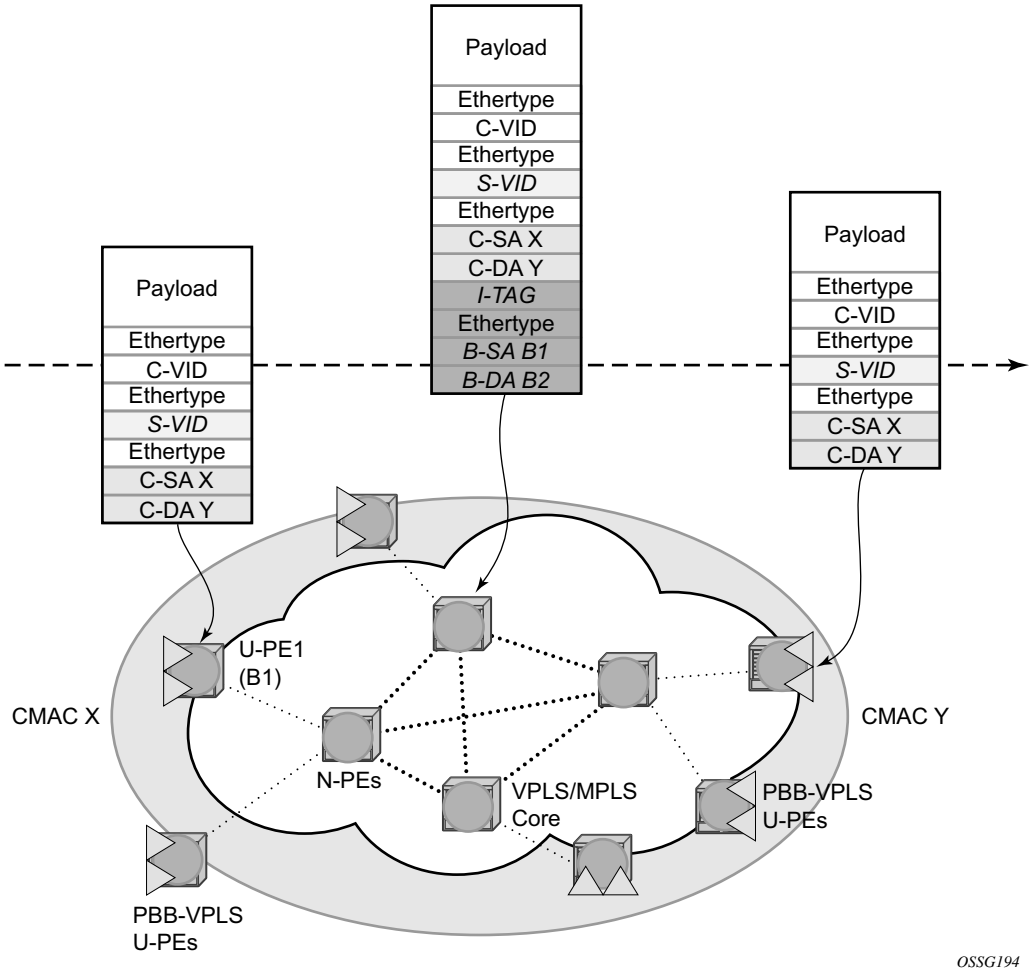

Figure 62: PBB Packet Walkthrough

The station with CMAC (customer MAC) X wants to send a unicast frame to CMAC Y through the PBB-VPLS network. A customer frame arriving at PBB-VPLS U-PE1 is encapsulated with the PBB header. The local I-VPLS FIB on U-PE1 is consulted to determine the destination BMAC of

the egress U-PE for CMAC Y. In our example, B2 is assumed to be known as the B-DA for Y. If CMAC Y is not present in the U-PE1 forwarding database, the PBB packet is sent in the B-VPLS using the standard group MAC address for the ISID associated with the customer VPN.

Next, only the Backbone Header in green is used to switch the frame through the green B-VPLS/ VPLS instances in the N-PEs. At the receiving U-PE2, the CMAC X is learned as being behind BMAC B1; then the PBB encapsulation is removed and the lookup for CMAC Y is performed.

PBB ELINE Service

ELINE service is defined in PBB (IEEE 802.1ah) as a point-to-point service over the B-component infrastructure. Alcatel-Lucent's implementation offers support for PBB ELINE through the mapping of multiple Epipe services to a Backbone VPLS infrastructure.

The use of Epipe scales the ELINE services as no MAC switching, learning or replication is required in order to deliver the point-to-point service.

All packets ingressing the customer SAP are PBB encapsulated and unicasted through the B-VPLS "tunnel" using the backbone destination MAC of the remote PBB PE.

All the packets ingressing the B-VPLS destined for the Epipe are PBB de-encapsulated and forwarded to the customer SAP.

PBB Resiliency for PBB epipe service

The PBB epipe service can be protected using G.8032 (the G8032 instance is created to protect the PBB B-VPLS service). For more information and for an example see Overview of G.8032 Operation.

PBB Resiliency for B-VPLS

The following VPLS resiliency mechanisms are also supported in PBB VPLS:

- Native Ethernet resiliency supported in both I-VPLS and B-VPLS contexts
- Distributed LAG, MC-LAG, RSTP
- MSTP in a management VPLS monitoring (B- or I-) SAPs.
- The G.8032 is supported for B-VPLS service. The G.8032 support is used only with PBB Epipe service from the current releases and cannot be used with PBB I-VPLS service.

Access Multi-Homing for Native PBB (B-VPLS over SAP Infrastructure)

Alcatel-Lucent PBB implementation allows the operator to use a native Ethernet infrastructure as the PBB core. Native Ethernet tunneling can be emulated using Ethernet SAPs to interconnect the related B-VPLS instances. This kind of solution might fit certain operational environments where Ethernet services was provided in the past using QinQ solution. The drawback is that no LDP signaling is available to provide support for Access Multi-homing for Epipe (pseudowire Active/ Standby status) or I-VPLS services (LDP MAC Withdraw). An alternate solution is required.

A PBB network using Native Ethernet core is depicted in Figure 63. MC-LAG is used to multi-home a number of edge switches running QinQ to PBB BEBs.

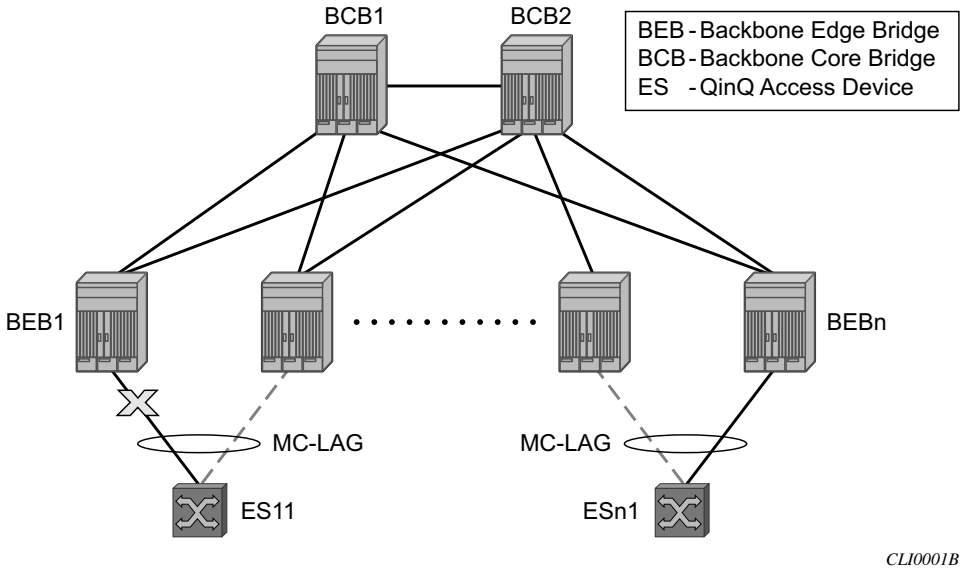

Figure 63: Access Dual-Homing into PBB BEBs - Topology View

The interrupted line from the MC-LAG represents the standby, inactive link; the solid line is the active link. The BEBs are dual-homed to two core switches BCB1 and BCB2 using native Ethernet SAPs on the B-VPLS side. Multi-point B-VPLS with MSTP for loop avoidance can be used as the PBB core tunneling.

PBB QoS

The following QoS processing rules apply for PBB B-VPLS SAPs:

B-VPLS SAP ingress

- If dot1p classification is enabled, the BTAG fields will be used by default to evaluate the internal forwarding class (fc) and discard profile if there is a BTAG field.
- If dot1p classification is not explicitly enabled or the packets are untagged then the default fc and profile is assigned.

B-VPLS SAP egress

- If the access port based policy contains FC and profile to dot1p mapping, this entry is used to mark the dot1p bits in the B-TAG of the frame going out of the SAP. The I-Tag of the frame is not modified in any case.
- If no explicit mapping exists, the related dot1p DE bits are set to zero on both ITAG and BTAG if the frame is originated locally from an I-VPLS. If the frame is transiting the B-VPLS the ITAG stays unchanged, the BTAG is set according to the type of ingress SAP.
  - → If the ingress SAP is tagged, the values of the dot1p, DE bits are preserved in the BTAG going out on the egress SAP.
  - → If the ingress SAP is untagged, the dot1p, DE bits are set to zero in the BTAG going out on the egress SAP.

I-SAP Ingress

- SAP ingress classification using mac-criteria or IP DSCP is supported.

I-SAP Egress

- Access port based marking is supported for I-SAPs (dot1q and QinQ SAPs).

PBB ACL Support

Filter policies are supported for ingress and egress of PBB I-SAP in both PBB epipe and PBB VPLS service.

Only MAC criteria Filter policies is available for use with PBB B-SAPs on ingress with the following functionality:

- For PBB B-VPLS B-SAPs, the MAC filter matches the outer MAC header fields (that is, B-DA, B-SA, B-Tag) for traffic received on a B-SAP and forwarded to another B-SAP in the system.
- For PBB B-VPLS B-SAPs, the MAC filter matches the inner MAC header fields (that is, the customer MAC DA, SA and VLAN tags) for traffic received on a B-SAP and forwarded out of an I-SAP in the system.

Only MAC criteria filter policies is available for use with PBB B-SAPs on egress. This filter policy only matches the BCB traffic. BEB traffic (that is, PBB originated traffic) cannot be matched using the egress filter policy attached to PBB B-SAP.

Configuration Guidelines

Listed below are the configuration guidelines for a PBB service:

- PBB services are supported only on 7210 SAS devices configured in network mode.
- A PBB service instance (identified by the ISID) cannot be used to encapsulate customer payloads with additional VLAN tags, if that service instance is being used to transport frames received on a QinQ access SAP. If a particular service instance is in use by a QinQ access SAP, then the system drops the packets that are received with additional tags on all the SAPs (NULL or Dot1q) using the same instance. Packets received with one or more tags on a NULL SAP, more than one tag on a Dot1q SAP, and more than two tags on a QinQ SAP are classified as packets with additional VLAN tags.
- Service MTU is not available for use.
- Port-based SHG is available for use with I-VPLS and B-VPLS service. Service based SHG is not available for use in an I-VPLS and a B-VPLS service.
- The system uses the internal loopback to flood/replicate BUM traffic received on the B-SAP, to create an additional copy for processing in the I-VPLS context. The system also uses the internal loopback to for egress port mirroring. The user needs to ensure that aggregate amount of mirrored traffic in the system and the BUM traffic received on a B-SAP does not exceed the available internal loopback bandwidth. Ingress meters can be used to limit the amount of BUM received and processed from a B-SAP and user can limit the number of ports setup for port egress mirroring to control the maximum amount of

traffic that needs to be circulated for two pass processing using the internal loopback. NOTE: If only PBB Epipe is used (no I-VPLS service is configured for use), then egress port mirroring can be enabled without affecting PBB traffic, since PBB Epipe traffic does not use the two-pass approach.

- Multiple B-SAPs on the same port cannot be part of the same B-VPLS service. Two B-SAPs on the same port need to be configured in two different services.
- Processing rules for packets received with multiple B-tags on a SAP:
  - → If the B-Tag header has two tags, the packet is processed and forwarded appropriately and sent out of an I-SID service or another B-VPLS B-SAP.
  - → If the node is acting as a pure BCB (with no ISID/service termination), then the packets are flooded and switched appropriately and if the node is acting as a BCB + BEB, then the packets are flooded and switched appropriately on the B-SAPs, but they will not switched or flooded to a I-SAPs (both VPLS and Epipe I-SAPs).
- PBB I-tag etype is not configurable, it is set to 0x88e7.
- PBB B-tag etype is not configurable; it is set to 0x8100.
- PBB packets received from a destination MAC address other than the one configured in the epipe service is not accepted by 7210 devices.
- In the current release, PBB packets with UCA bit set are dropped.
- Aging of MAC addresses learnt in the B-domain - As long as a Customer MAC (C-MAC) or an Epipe service is associated with an B-SA/B-MAC, do not age out the B-SA. When the last customer MAC ages out or the last epipe service using the particular B-SA MAC is removed, remove the corresponding B-SA entry. This means that as long as an epipe service is associated with a particular PBB destination MAC address, the corresponding B-MAC will not age out and will occupy an entry in the L2 learning table. Note, that if only I-VPLS is in use, then aging out of C-MAC will automatically trigger aging out B-MAC, when the last C-MAC associated with the B-MAC is aged out.

Configuration Guidelines (for 7210 SAS-M and 7210 SAS-T)

Listed below are the configuration guidelines specific to 7210 SAS-M and 7210 SAS-T devices configured in Network mode:

When “discard-unknown” is enabled on a B-VPLS, the following behavior can be observed:

- Unknown unicast (B-DA) packets arriving on a B-SAP are dropped.
- Unknown unicast (C-DA) packets arriving on a B-SAP are processed normally in the I-VPLS, if the B-DA is not unknown unicast.
- Unknown unicast (C-DA) packets arriving on an I-SAP are not dropped and are flooded in the B-VPLS, because B-DA is equal to the “Group Mcast MAC” and is a known value

- Mac-protect feature is not available for use in I-VPLS or B-VPLS service
- Port based SHG is available for use with both I-VPLS and B-VPLS service. Service based SHG is not available in both.

Configuration Examples

Use the CLI syntax displayed to configure PBB.

PBB ELAN and ELINE

Use the following CLI syntax to bring up PBB B-VPLS - common to both ELAN and ELINE services:

CLI Syntax:
```
config>service# vpls 200 customer 1 b-vpls create
            description “This is a B-VPLS.”
            sap 3/1/3:33 create
               description “B-VPLS SAP”
```

Use the following CLI syntax to bring up PBB ELAN:

CLI Syntax:
```
config>service# vpls 2000 customer 6 i-vpls create
            description “This is an I-VPLS.”
            sap 4/1/3:20 create
               description “I-VPLS SAP”
            backbone-vpls 200
```

Use the following CLI syntax to bring up PBB ELINE:

CLI Syntax:
```
config>service# epipe 1000 customer 10 create pbb-epipe
            description “This is an Epipe.”
            sap 4/1/3:20 create
               description “Epipe SAP”
            pbb-tunnel 200 backbone-dest-mac 00-01-10-1E-C6-67 isid 752
```

MC-LAG Multihoming for Native PBB

This section describes a configuration example for BEB C configuration given the following assumptions:

- BEB C and BEB D are MC-LAG peers
- B-VPLS 100 on BEB C and BEB D
- VPLS 1000 on BEB C and BEB D
- MC-LAG 1 on BEB C and BEB D

CLI Syntax:

```
service pbb
      source-bmac ab-ac-ad-ef-00-00
port 1/1/1
      ethernet
            encap-type qinq
lag 1
      port 1/1/1 priority 20
      lacp active administrative-key 32768
redundancy
      multi-chassis
            peer 1.1.1.3 create
                  source-address 1.1.1.1
                  mc-lag
                        lag 1 lacp-key 1 system-id 00:00:00:01:01:01
                        system-priority 100
                              source-bmac-lsb use-lacp-key
service vpls 100 bvpls
      sap 2/2/2:100 // bvid 100
      mac-notification
            no shutdown

service vpls 101 bvpls
      sap 2/2/2:101 // bvid 101
      mac-notification
            no shutdown
// no per BVPLS source-bmac configuration, the chassis one (ab-ac-ad-ef-
00-00) is used

service vpls 1000 ivpls
      backbone-vpls 100
      sap lag-1:1000 //automatically associates the SAP with ab-ac-ad-
      ef-00-01 (first 36 bits from BVPLS 100 sbmac+16bit source-bmac-
      lsb)
```

```
service vpls 1001 ivpls
      backbone-vpls 101
      sap lag-1:1001 //automatically associates the SAP with ab-ac-ad-
      ef-00-01(first 36 bits from BVPLS 101 sbmac+16bit source-bmac-lsb)
```

PBB Command Reference

Command Hierarchies


```
config
    — service
        — pbb
            — mac-name name ieee-address
            — no mac-name
            — source-bmac ieee-address
            — no source-bmac

config
    — service
        — [no] vpls service-id [customer customer-id] [create] ] [vpn vpn-id] [m-vpls] [svc- sap-type
          {null-star|dot1q-preserve|any}] [customer-vid vlan-id] [b-vpls] [i-vpls]
            — backbone-vpls service-id [isid isid]
            — no backbone-vpls
            — pbb
                — no source-bmac ieee-address

config
    — service
        — [no] epipe service-id [customer customer-id] [create] [vpn vpn-id] svc-sap-type {null-
          star|dot1q-preserve|any}] [customer-vid vlan-id] [pbb-epipe]
            — [no] pbb-tunnel service-id backbone-dest-mac mac-name| ieee-address isid ISID
```

Show Commands

```
show
    — eth-cfm
        — association [ma-index] [detail]
        — cfm-stack-table [port port-id [vlan qtag[.qtag]] | sdp sdp-id[:vc-id] [level 0..7] [direction
          up|down]
        — domain [md-index] [association ma-index | all-associations [detail]]
        — mep mep-id domain md-index association ma-index [loopback] [linktrace]
        — mip
    — service
        — id service-id
            — i-vpls
            — epipe
```

— **all**
— **base**
— **fdb {info | mac ieee-address | sap sap-id | detail | endpoint endpoint} [expiry] [pbb]**
— **stp [detail]**
— **isid-using [***ISID***]**
— **pbb**
— **base**
— **mac-name** [**detail**]
— **mac-name**
— **mac-name** *mac-name* **detail**
— **service-using [b-vpls] [i-vpls]**

Clear Commands

clear
— **service**
— **id** *service-id*
— **fdb {all | mac** *ieee-address* | *sap sapid***}**
— **stp**
— **detected-protocols [all | sap** *sap-id***]**
— **statistics**
— **id** *service-id*
— **counters**
— **stp**
— **sap** *sap-id* {**all | counters | stp** }

Debug Commands

debug
— **service**
— **id** *service-id*
— **[no] event-type {config-change | svc-oper-status-change | sap-oper-status-change}**|
— **[no] sap** *sap-id*
— **stp**
— **all-events**
— **[no] bpdu**
— **[no] core-connectivity**
— **[no] exception**
— **[no] fsm-state-changes**
— **[no] fsm-timers**
— **[no] port-role**
— **[no] port-state**
— **[no] sap** *sap-id*

PBB Service Commands

VPLS Service Commands

vpls

Syntax **vpls** *service-id* [**customer** *customer-id*] [**create**][**vpn** *vpn-id*] [*m-vpls*]
service-id [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [m-vpls] [**svc-sap-type** {**null-star|dot1q-preserve|any**}] [**customer-vid** *vlan-id*] [*b-vpls*]
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.

Parameters **any** — Allows any SAP type. When svc-sap-type is set to any, for a NULL SAP, the system processes and forwards only packets with no VLAN tag (that is, untagged). All other packets with one or more VLAN tags (even those with priority tag only) are not processed and dropped. Users can use the service with svc- sap-type set to `null-star' to process and forward packets with one or more tags (including priority tag) on a null SAP.

Default null-star

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 7210 SAS on which this service is defined.

Values 1 — 2147483648

b-vpls — Creates a backbone-vpls.

create — This keyword is mandatory while creating a VPLS service.

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

customer-vid vlan-id — Defines the dot1q VLAN ID to be specified while creating the local Dot1q SAP for svc-sap-type dot1q-preserve.

Default 1 — 4094

dot1q-preserve — Specifies that the allowed SAP in the service are Dot1q. The Dot1q ID is not stripped after packets matches the SAP.

Default null-star

m-vpls — Specifies a management VPLS.

null-star — Specifies that the allowed SAP in the service are either null SAPs or Dot1q* SAPs.

svc-sap-type — Specifies the type of service and allowed SAPs in the service.

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)

pbb

Syntax **pbb**

Context config>service
config>service>vpls

Description This command configures PBB parameters.

mac-name

Syntax **mac-name** *name ieee-address*
no mac-name *name*

Context config>service>pbb

Description This command configures the MAC name for the MAC address. It associates an ASCII name with an IEEE MAC to improve the PBB Epipe configuration. It can also change the dest-BMAC in one place instead of 1000s of Epipe.

Parameters *name* — Specifies the MAC name up to 32 characters in length.

ieee-address — The MAC address assigned to the MAC name. The value should be input in either a xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format.

source-bmac

Syntax **source-bmac** *ieee-address*
no source-bmac

Context config>service>pbb

Description This command configures the base source BMAC for the B-VPLS. The first 32 bits must be the same with what is configured in the MC-LAG peer. If the base source BMAC is not configured, it inherits the chassis level BMAC configured under the PBB object added in the previous section.

Parameters *ieee-address* — The MAC address assigned to the BMAC. The value should be input in either a xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format.

backbone-vpls

Syntax **backbone-vpls** *service-id* [**isid** *isid*]
no backbone-vpls

Context config>service>vpls
config>service>vpls>pbb

Description This command configures B-VPLS service associated with the I-VPLS.

Parameters *service-id* — Specifies the service ID.

Values 1..2147483648

isid — Specifies the ISID.

Values 0..16777215

Epipe Service Commands

epipe

Syntax **epipe** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*][**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**svc-sap-type {null-star|dot1q|dot1q-preserve|any}**] [**customer-vid** *vlan-id*] [**pbb-epipe**]
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 7210 SAS.

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.

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 7210 on which this service is defined.

Values 1 — 2147483648

any — When svc-sap-type is set to any, for a NULL SAP, the system processes and forwards only packets with no VLAN tag (that is, untagged). All other packets with one or more VLAN tags (even those with priority tag only) are not processed and dropped. Users can use the service with svc- sap-type set to null-star, to process and forward packets with one or more tags (including priority tag) on a null SAP.

Default null-star

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

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

customer-vid vlan-id — Defines the dot1q VLAN ID to be specified while creating the local Dot1q SAP for svc-sap-type dot1q-preserve.

Values 1 — 4094

dot1q — Specifies that the allowed SAP in the service are Dot1q SAPs and dot1q explicit null SAPs.

dot1q-preserve — Specifies that the allowed SAP in the service are Dot1q. The Dot1q ID is not stripped after packet matches the SAP.

null-star — Specifies that the allowed SAP in the service are either null SAPs or Dot1q default SAPs.

pbb-epipe — keyword used to create a pbb-epipe.

svc-sap-type — Specifies the type of service and allowed SAPs in the service.

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)

pbb-tunnel

Syntax **pbb-tunnel** *service-id* **backbone-dest-mac** {*mac-name* | *ieee-mac*} **isid** *ISID*
no pbb-tunnel

Context config>service>epipe

Description This command configures a Provider Backbone Bridging (PBB) tunnel with Backbone VPLS (B-VPLS) service information.

Parameters *service-id* — Specifies the B-VPLS service for the PBB tunnel associated with this service.

Values 1 — 2147483648

backbone-dest-mac {*mac-name* | *ieee-mac*} **—** Specifies the backbone destination MAC-address for PBB packets.

isid *ISID* **—** Specifies a 24 bit service instance identifier for the PBB tunnel associated with this service. As part of the PBB frames, it is used at the destination PE as a demultiplexor field.

Values 0 — 16777215

PBB Show Commands

eth-cfm

Syntax **eth-cfm**

Context show

Description This command displays 802.1ag CFM information.

association

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

Context show>eth-cfm

Description Shows association information.

Parameters *ma-index* — Specifies the MA index value.

Values 1 — 4294967295

detail — Displays all association detail.

Output

```
*A:alcag1-R6# show eth-cfm association
===============================================================================
CFM Association Table
===============================================================================
Md-index   Ma-index   Name                      CCM-interval Bridge-id
-------------------------------------------------------------------------------
1          1          ivpls                     1            5000
===============================================================================
*A:alcag1-R6#
```

cfm-stack-table

Syntax **cfm-stack-table**
cfm-stack-table port [*port-id*> [**vlan** qtag[.qtag]] [**level** *0..7*] [**direction up | down**]
cfm-stack-table sdp [*sdp-id[:vc-id]*>] [**level** *0..7*]] [**direction up | down**]
cfm-stack-table virtual [*service-id*] [**level** *0..7*]

Context show>eth-cfm

Description Summarizes all MEPs/MIPs.

Parameters *port-id* — Displays information about the specified port.

Values port-id slot/mda/port[.channel]
lag-id lag-*id*

lag	keyword
id	1 — 200

sdp-id[:*vc-id*] — Specfies an existing SDP and VC ID.

Values 1 — 17407

qtag — Specifies the qtag value.

Values 0 — 4094

level — Specifies the level.

Values 0 — 7

direction up | down — Indicates the direction in which the maintenance association (MEP or MIP) faces on the bridge port.

down — Displays continuity check information configured away from the MAC relay entity.

up — Displays continuity check information configured toward the MAC relay entity.

service-id — Specifies information about the specified service ID.

Values 1 — 2147483648

Sample Output

```
*A:alcag1-R6# show eth-cfm cfm-stack-table
===============================================================================
CFM SAP Stack Table
===============================================================================
Sap            Level Dir  Md-index   Ma-index   Mep-id Mac-address
-------------------------------------------------------------------------------
1/2/9:5        4     Up   1          1          51     00:ae:ae:ae:ae:ae
===============================================================================
CFM SDP Stack Table
===============================================================================
Sdp            Level Dir  Md-index   Ma-index   Mep-id Mac-address
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:alcag1-R6#
```

domain

Syntax **domain** [*md-index*] [**association** *ma-index* | **all-associations** [**detail**]]

Context show>eth-cfm>domain

Description This command displays domain information.

Parameters *md-index* — Specifies the maintenance domain (MD) index value.

Values 1 — 4294967295

ma-index — Specifies the MA index value.

Values 1 — 4294967295

all-associations — Displays information all maintenance associations.

detail — Displays detailed information.

Sample Output

```
*A:alcag1-R6# show eth-cfm domain
===============================================================================
CFM Domain Table
===============================================================================
Md-index   Level Name                                          Format
-------------------------------------------------------------------------------
1          4     ivpls                                        charString
===============================================================================
*A:alcag1-R6#

*A:alcag1-R6# show eth-cfm mep 51 domain 1 association 1
-------------------------------------------------------------------------------
Mep Information
-------------------------------------------------------------------------------
Md-index           : 1                        Direction         : Up
Ma-index           : 1                        Admin             : Enabled
MepId              : 51                       CCM-Enable        : Enabled
IfIndex            : 38043648                 PrimaryVid        : 5
FngState           : fngReset
LowestDefectPri    : allDef                   HighestDefect     : none
Defect Flags       : None
Mac Address        : 00:ae:ae:ae:ae:ae        CcmLtmPriority    : 7
CcmTx              : 775                      CcmSequenceErr    : 0
CcmLastFailure Frame:
    None
XconCcmFailure Frame:
    None
*A:alcag1-R6#
```

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**loopback**] [**linktrace**]

Context show>eth-cfm>domain

Description This command displays Maintenance Endpoint (MEP) information.

Parameters *mep-id* — Specifies the maintenance association end point identifier.

Values 1 — 8191

md-index — Specifies the maintenance domain (MD) index value.

Values 1 — 4294967295

ma-index — Specifies the MA index value.

Values 1 — 4294967295

loopback — Displays loopback information for the specified MEP.

linktrace — Displays linktrace information for specified MEP.

Sample Output

```
*A:alcag1-R6# oam eth-cfm loopback 00:af:af:af:af:af mep 51 domain 1 association 1
eth-cfm Loopback Test Initiated: Mac-Address: 00:af:af:af:af:af, out sap: 1/2/9:5
Sent 1 packets, received 1 packets [0 out-of-order, 0 Bad Msdu] -- OK
*A:alcag1-R6#


*A:alcag1-R6# oam eth-cfm linktrace 00:af:af:af:af:af mep 51 domain 1 association 1
Index Ingress Mac          Egress Mac           Relay      Action
----- -------------------- -------------------- ---------- ----------
1     00:00:00:00:00:00    00:AF:AF:AF:AF:AF    rlyHit     terminate
----- -------------------- -------------------- ---------- ----------
No more responses received in the last 5 seconds.
*A:alcag1-R6#
```

mip

Syntax **mip**

Context show>eth-cfm>mip

Description This command displays Maintenance Intermediate Point (MIP) information only in 7210 SAS-M and 7210 SAS-T network mode and SASX.

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

Label	Description
Mip-Enabled	Displays the state of the MIP service
Mip Mac Address	Indicates the Mac Address of the MIP

Sample

```
*A:7210SAS#  show eth-cfm mip
===============================================================================
CFM SAP MIP Table
===============================================================================
Sap                                   Mip-Enabled     Mip Mac Address
-------------------------------------------------------------------------------
1/1/16                                yes             00:a1:b1:c1:d1:e1
===============================================================================
===============================================================================
CFM SDP MIP Table
===============================================================================
Sdp                                   Mip-Enabled     Mip Mac Address
-------------------------------------------------------------------------------
456:123                               yes             00:a2:b2:c2:d2:e2
===============================================================================
*A:7210SAS#
```

id

Syntax **id** *service-id*

Context show>service

Description This command displays information on a specific service ID.

Parameters *service-id* — The unique service identification number that identifies the service in the service domain.

Values service-id: 1 — 214748364

all — Displays detailed information about the service.

base — Displays basic service information.

fdb — Displays FDB entries.

epipe — Displays the e-pipe services associated with the B-VPLS service.

i-vpls — Displays the I-VPLS services associated with this B-VPLS service.

stp — Display STP information.

all

Syntax **all**

Context show>service>id

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 Id	The service identifier.
Service Type	Specifies the type of service.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Status Change	The date and time of the most recent status change to this customer.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.

Label	Description
Admin State	The administrative state of the service
Vc Switching	Displays the status of VC switching.
SAP Count	The number of SAPs specified for this service.
Optim	Displays the mode of the device.
Vpn Id	The number which identifies the VPN.
Oper State	The operational state of the service.
SAP	Displays the SAP ID.
Encap	Indicates the encapsulation value used to identify the SAP on the access port.
Qinq Ethertype	Displays the configured Q-in-Q Ethertype value.
Dot1Q Ethertype	Displays the configured Dot1Q Ethertype value
Ingress Qos-policy	Displays the ingress QoS policy ID.
Egress Qos-policy	Displays the egress QoS policy ID.
Admin MTU	Displays the desired largest service frame size (in octets) that can be transmitted through the port to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the port to the far-end router, without requiring the packet to be fragmented.
Ingress IP Fltr-Id	The ingress IP filter ID.
Egress IP Fltr-Id	The egress IP filter ID.
Ingress MAC Fltr-Id	The ingress MAC filter ID.
Egress MAC Fltr-Id	The egress MAC filter ID.
Ingress IPv6 Fltr-Id	The ingress IPv6 filter ID.
Egress IPv6 Fltr-Id	The egress IPv6 filter ID.
SAP	Displays the SAP ID.
Acct. Pol	Displays the accounting policy ID assigned to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
QoS parameters	
Ingress qos-policy	Displays the ingress QoS policy ID.
Egress qos-policy	Displays the egress QoS policy ID.
Classifiers Used	Displays the number of classifiers used.

Label	Description
Ingr IP Fltr-Id	Ingress IP filter ID.
Egr IP Fltr-Id	Egress IP filter ID
Ingr Mac Fltr-Id	Ingress MAC filter ID
Egr Mac Fltr-Id	Egress MAC filter ID
Ingr IPv6 Fltr-Id	Ingress IPv6 filter ID
Egr IPv6 Fltr-Id	Egress IPv6 filter ID
Endpoint	Displays the endpoint name
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
	QoS parameters
Ingress qos-policy	The SAP ingress QoS policy ID.
Classifiers Allocated	Displays the number of classifiers allocated.
Classifiers Used	Displays the number of classifiers used.

Sample Output

```
Sample output for PBB Epipe:

*A:7210-SAS>show>service# id 1000 all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 1000               Vpn Id            : 0
Service Type      : Epipe
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 04/04/2001 22:18:48
Last Mgmt Change  : 04/04/2001 21:28:34
Admin State       : Up                 Oper State        : Up
MTU               : n/a
MTU Check         : n/a
Vc Switching      : False
SAP Count         : 1                  SDP Bind Count    : 0
Uplink Type:      : MPLS
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/1/15:1000
-------------------------------------------------------------------------------
Service Id         : 1000
SAP                : 1/1/15:1000             Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Up
Flags              : None
Last Status Change : 04/04/2001 21:29:23
Last Mgmt Change   : 04/04/2001 21:28:34
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU          : 1518                    Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
tod-suite          : None
Endpoint           : N/A

Acct. Pol          : None                    Collect Stats     : Disabled


-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1
-------------------------------------------------------------------------------
Aggregate Policer
-------------------------------------------------------------------------------
rate               : n/a                     burst             : n/a
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
```

```
Classifiers Allocated: 4                    Meters Allocated   : 2
Classifiers Used     : 1                    Meters Used        : 1
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress Stats:          0                       0
Egress Stats:           0                       0
Ingress Drop Stats:     0                       0

Extra-Tag Drop Stats:   n/a                     n/a
-------------------------------------------------------------------------------
Sap per Meter stats (in/out counter mode)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Meter 1 (Unicast)
For. InProf           : 0                       0
For. OutProf          : 0                       0

-------------------------------------------------------------------------------
PBB Tunnel Point
-------------------------------------------------------------------------------
B-vpls     Backbone-dest-MAC Isid     AdmMTU OperState Flood Oper-dest-MAC
-------------------------------------------------------------------------------
2          8c:90:d3:79:b2:65 1000     1514   Up        Yes   8c:90:d3:79:b2:65
-------------------------------------------------------------------------------
Last Status Change: 04/04/2001 22:18:48
Last Mgmt Change:   04/04/2001 22:18:48

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:7210-SAS>show>service#
```

Sample output for I-VPLS:

```
*A:7210-SAS>show>service# id 200 all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 200                 Vpn Id            : 0
Service Type      : i-VPLS
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 04/04/2001 22:14:30
Last Mgmt Change  : 04/04/2001 22:15:06
Admin State       : Up                  Oper State        : Up
MTU               : n/a
MTU Check         : n/a
SAP Count         : 1                   SDP Bind Count    : 0
Snd Flush on Fail : Disabled
Uplink Type:      : MPLS
b-Vpls Id         : 2                   Oper ISID         : 200
b-Vpls Status     : Up
```

```
-------------------------------------------------------------------------------
Split Horizon Group specifics
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/1/15:200
-------------------------------------------------------------------------------
Service Id         : 200
SAP                : 1/1/15:200              Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Up
Flags              : None
Last Status Change : 04/04/2001 22:14:30
Last Mgmt Change   : 04/04/2001 22:14:22
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
Admin MTU          : 1518                    Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
tod-suite          : None
Mac Learning       : Enabled                 Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                 Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled

Acct. Pol          : None                    Collect Stats     : Disabled


-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                      Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Forwarding
Port Number        : 2049                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : N/A
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                     Designated Port   : N/A

Forward transitions: 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
TC bit BPDUs rcvd  : 0                       TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                       MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
```

```
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                        Min Auth Interval : 15 minutes

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1
-------------------------------------------------------------------------------
Aggregate Policer
-------------------------------------------------------------------------------
rate               : n/a                      burst             : n/a
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
Classifiers Allocated: 4                    Meters Allocated   : 2
Classifiers Used     : 2                    Meters Used        : 2
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                       Packets                Octets
Ingress Stats:         0                      0
Egress Stats:          0                      0
Ingress Drop Stats:    0                      0

Extra-Tag Drop Stats:  n/a                    n/a
-------------------------------------------------------------------------------
Sap per Meter stats (in/out counter mode)
-------------------------------------------------------------------------------
                       Packets                Octets

Ingress Meter 1 (Unicast)
For. InProf          : 0                      0
For. OutProf         : 0                      0

Ingress Meter 11 (Multipoint)
For. InProf          : 0                      0
For. OutProf         : 0                      0

-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up                     Core Connectivity : Down
Stp Admin State    : Down                   Stp Oper State    : Down
Mode               : Rstp                   Vcp Active Prot.  : N/A

Bridge Id          : 80:00.00:25:ba:08:f6:20  Bridge Instance Id: 0
Bridge Priority    : 32768                  Tx Hold Count     : 6
Topology Change    : Inactive               Bridge Hello Time : 2
Last Top. Change   : 0d 00:00:00            Bridge Max Age    : 20
Top. Change Count  : 0                      Bridge Fwd Delay  : 15

Root Bridge        : N/A
Primary Bridge     : N/A

Root Path Cost     : 0                      Root Forward Delay: 0
Rcvd Hello Time    : 0                      Root Max Age      : 0
Root Priority      : 0                      Root Port         : N/A

-------------------------------------------------------------------------------
Forwarding Database specifics
```

```
-------------------------------------------------------------------------------
Service Id        : 200                Mac Move          : Disabled
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                Total Count       : 0
Learned Count     : 0                  Static Count      : 0
Remote Age        : 900                Local Age         : 300
High Watermark    : 95%                Low Watermark     : 90%
Mac Learning      : Enabled            Discard Unknown   : Disabled
Mac Aging         : Enabled            Relearn Only      : False
===============================================================================
```

Sample output for B-VPLS service:

```
*A:7210-SAS>show>service# id 2 all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 2                  Vpn Id            : 0
Service Type      : b-VPLS
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 04/04/2001 22:13:57
Last Mgmt Change  : 04/04/2001 22:13:57
Admin State       : Up                 Oper State        : Up
MTU               : n/a
MTU Check         : n/a
SAP Count         : 1                  SDP Bind Count    : 0
Snd Flush on Fail : Disabled
Uplink Type:      : MPLS
Oper Backbone Src : 00:25:ba:08:f6:20
i-Vpls Count      : 1
Epipe Count       : 1

-------------------------------------------------------------------------------
Split Horizon Group specifics
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/1/2:2
-------------------------------------------------------------------------------
Service Id         : 2
SAP                : 1/1/2:2                Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                     Oper State        : Up
Flags              : None
Last Status Change : 04/04/2001 22:13:57
Last Mgmt Change   : 04/04/2001 22:13:54
Dot1Q Ethertype    : 0x8100                 QinQ Ethertype    : 0x8100
PBB Ethertype      : 0x88e7
```

```
Split Horizon Group: (Not Specified)

Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
Admin MTU          : 1518                    Oper MTU          : 1518
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
tod-suite          : None
Mac Learning       : Enabled                 Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                 Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled

Acct. Pol          : None                    Collect Stats     : Disabled


-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                      Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Forwarding
Port Number        : 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
TC bit BPDUs rcvd  : 0                       TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                       MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                       Min Auth Interval : 15 minutes

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1
-------------------------------------------------------------------------------
Aggregate Policer
-------------------------------------------------------------------------------
rate               : n/a                     burst             : n/a
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
Classifiers Allocated: 4                    Meters Allocated   : 2
Classifiers Used     : 2                    Meters Used        : 2
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                      Packets                 Octets
Ingress Stats:        0                       0
Egress Stats:         0                       0
Ingress Drop Stats:   0                       0
```

```
Extra-Tag Drop Stats:   n/a                     n/a
-------------------------------------------------------------------------------
Sap per Meter stats (in/out counter mode)
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Meter 1 (Unicast)
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Meter 11 (Multipoint)
For. InProf           : 0                       0
For. OutProf          : 0                       0

-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up                       Core Connectivity : Down
Stp Admin State    : Down                     Stp Oper State    : Down
Mode               : Rstp                     Vcp Active Prot.  : N/A

Bridge Id          : 80:00.00:25:ba:08:f6:20  Bridge Instance Id: 0
Bridge Priority    : 32768                    Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 0d 00:00:00              Bridge Max Age    : 20
Top. Change Count  : 0                        Bridge Fwd Delay  : 15

Root Bridge        : N/A
Primary Bridge     : N/A

Root Path Cost     : 0                        Root Forward Delay: 0
Rcvd Hello Time    : 0                        Root Max Age      : 0
Root Priority      : 0                        Root Port         : N/A

-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id        : 2                  Mac Move          : Disabled
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                Total Count       : 0
Learned Count     : 0                  Static Count      : 0
Remote Age        : 900                Local Age         : 300
High Watermark    : 95%                Low Watermark     : 90%
Mac Learning      : Enabled            Discard Unknown   : Disabled
Mac Aging         : Enabled            Relearn Only      : False

-------------------------------------------------------------------------------
Related i-Vpls services for b-Vpls service 2
-------------------------------------------------------------------------------
i-Vpls SvcId        Oper ISID           Admin               Oper
-------------------------------------------------------------------------------
200                 200                 Up                  Up
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Related Epipe services for b-Vpls service 2
-------------------------------------------------------------------------------
```

```
Epipe SvcId         Oper ISID           Admin               Oper
-------------------------------------------------------------------------------
1000                1000                Up                  Up
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:7210-SAS>show>service#
```

base

Syntax **base**

Context show>service>id

Description This command displays basic information about the service including service type, description and SAPs.

Output Show service ID base output — The following table describes the command output fields.

Label	Description
Service Id	The service identifier.
Service Type	Specifies the type of service.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Status Change	The date and time of the most recent status change to this customer.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Admin State	The administrative state of the service
Vc Switching	Displays the status of VC switching .
SAP Count	The number of SAPs specified for this service.
Uplink Type	Displays the mode of the device.
Vpn Id	The number which identifies the VPN.
Oper State	The operational state of the service.

SAP	Displays the SAP ID.
Encap	The value of the label used to identify this SAP on the access port.
Vpn Id	The number which identifies the VPN.
Oper State	The operational state of the service.
SAP	Displays the SAP ID.
	PBB Tunnel Point
B-vpls	Displays the B-VPLS ID.
Backbone-dest-MAC	Displays the back bone destination MAC address.
Isid	Displays the ISID number.
Flood	Specifies whether or not the traffic is flooded in the B-VPLS for the Destination instead of unicast. If the backbone destination MAC is in theB-VPLS FDB, then it will be unicast.
b-Vpls Status	Displays the operational state of the B-VPLS service
b-Vpls Id	Displays the B-VPLS ID.

Sample

```
Sample output for PBB Epipe service:


*A:7210-SAS>show>service# id 1000 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1000               Vpn Id            : 0
Service Type      : Epipe
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 04/04/2001 22:18:48
Last Mgmt Change  : 04/04/2001 21:28:34
Admin State       : Up                 Oper State        : Up
MTU               : n/a
MTU Check         : n/a
Vc Switching      : False
SAP Count         : 1                  SDP Bind Count    : 0
Uplink Type:      : MPLS

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                            Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/15:1000                        q-tag        1518    1518    Up   Up
```

```
-------------------------------------------------------------------------------
PBB Tunnel Point
-------------------------------------------------------------------------------
B-vpls     Backbone-dest-MAC Isid     AdmMTU OperState Flood Oper-dest-MAC
-------------------------------------------------------------------------------
2          8c:90:d3:79:b2:65 1000     1514   Up        Yes   8c:90:d3:79:b2:65
-------------------------------------------------------------------------------
Last Status Change: 04/04/2001 22:18:48
Last Mgmt Change:   04/04/2001 22:18:48
===============================================================================
*A:7210-SAS>show>service#
```

Sample output for I-VPLS service:

```
*A:7210-SAS>show>service# id 200 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 200                 Vpn Id            : 0
Service Type      : i-VPLS
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 04/04/2001 22:14:30
Last Mgmt Change  : 04/04/2001 22:15:06
Admin State       : Up                  Oper State        : Up
MTU               : n/a
MTU Check         : n/a
SAP Count         : 1                   SDP Bind Count    : 0
Snd Flush on Fail : Disabled
Uplink Type:      : MPLS
b-Vpls Id         : 2                   Oper ISID         : 200
b-Vpls Status     : Up

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/15:200                           q-tag        1518    1518    Up   Up
===============================================================================
*A:7210-SAS>show>service#
```

Sample output for B-VPLS service:

```
*A:7210-SAS>show>service# id 2 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 2                   Vpn Id            : 0
Service Type      : b-VPLS
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 04/04/2001 22:13:57
Last Mgmt Change  : 04/04/2001 22:13:57
Admin State       : Up                  Oper State        : Up
```

```
MTU               : n/a
MTU Check         : n/a
SAP Count         : 1                   SDP Bind Count    : 0
Snd Flush on Fail : Disabled
Uplink Type:      : MPLS
Oper Backbone Src : 00:25:ba:08:f6:20
i-Vpls Count      : 1
Epipe Count       : 1

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/2:2                               q-tag        1518    1518    Up   Up
===============================================================================
*A:7210-SAS>show>service#
```

fdb

Syntax **fdb {info | mac** *ieee-address* **| sap** *sap-id* **| detail | endpoint** *endpoint***} [expiry] [pbb]**

Context show>service>id

Description This command displays FDB entries for a given MAC address.

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

detail — Displays detailed information.

expiry — Displays time until MAC is aged out.

endpoint — Displays endpoint information.

pbb — Displays PBB information.

Output Show FDB Information — The following table describes service FDB output fields:

Label	Description
Service Id	Displays the service ID.
Mac Move Rate	Displays the maximum rate at which MAC's can be re-learned in this 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 MAs. The rate is computed as the maximum number of re-learns allowed in a 5 second interval: for example, the default rate of 2 re-learns per second corresponds to 10 re-learns in a 5 second period.

Mac Move Retries	Displays the number of times retries are performed for re-enabling the SAP.
Table Size	Specifies the maximum number of learned and static entries allowed in the FDB of this service.
Learned Count	Displays the current number of learned entries in the FDB of this service.
Remote Age	Displays the number of seconds used to age out FDB entries learned on an SAP. These entries correspond to MAC addresses learned on remote SAPs.
High Watermark	Displays the utilization of the FDB table of this service at which a table full alarm will be raised by the agent.
Mac Learning	Specifies whether the MAC learning process is enabled
Mac Aging	Indicates whether the MAC aging process is enabled.
Mac Move	Displays the administrative state of the MAC movement feature associated with this service.
Mac Move Timeout	Displays the time in seconds to wait before a SAP that has been disabled after exceeding the maximum 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.
Total Count	Displays the total number of learned entries in the FDB of this service.
Static Count	Displays the current number of static entries in the FDB of this service.
Local Age	Displays the number of seconds used to age out FDB entries learned on local SAPs.
Low Watermark	Displays the utilization of the FDB table of this service at which a table full alarm will be cleared by the agent.
Discard Unknown	Specifies whether frames received with an unknown destination MAC are discarded.
Relearn Only	Displays, that when enabled, either the FDB table of this service is full, or that the maximum system-wide number of MAC's supported by the agent has been reached, and thus MAC learning is temporary disabled, and only MAC re-learns can take place.

Sample Output

```
*A:7210-SAS>show>service# id 200 fdb

===============================================================================
Forwarding Database, Service 200
===============================================================================
Service Id        : 200                 Mac Move          : Disabled
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                 Total Count       : 0
Learned Count     : 0                   Static Count      : 0
Remote Age        : 900                 Local Age         : 300
High Watermark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabled             Discard Unknown   : Disabled
Mac Aging         : Enabled             Relearn Only      : False
===============================================================================
*A:7210-SAS>show>service#

*A:7210-SAS>show>service# id 2 fdb

===============================================================================
Forwarding Database, Service 2
===============================================================================
Service Id        : 2                   Mac Move          : Disabled
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                 Total Count       : 0
Learned Count     : 0                   Static Count      : 0
Remote Age        : 900                 Local Age         : 300
High Watermark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabled             Discard Unknown   : Disabled
Mac Aging         : Enabled             Relearn Only      : False
===============================================================================
*A:7210-SAS>show>service#
```

stp

Syntax **stp [detail]**

Context show>service>id

Description This command 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
Bridge Id	Specifies the MAC address used to identify this bridge in the network.
Top. Change Count	Specifies the total number of topology changes detected by the Spanning Tree Protocol instance associated with this service since the management Entity was last reset or initialized.
Root Bridge	Specifies the bridge identifier of the root of the spanning tree as determined by the Spanning Tree Protocol instance associated with this service. This value is used as the Root Identifier parameter in all Configuration BPDUs originated by this node.
Stp Oper State	Displays the operational state of the STP
Primary Bridge	Specifies the bridge identifier of the root of the spanning tree as determined by the Spanning Tree Protocol instance associated with this service. This value is used as the Root Identifier parameter in all Configuration BPDUs originated by this node.
Topology Change	Specifies whether a topology change is currently in progress.
Mode	Displays the mode of the STP
Last Top. Change	Specifies the time (in hundredths of a second) since the last time a topology change was detected by the Spanning Tree Protocol instance associated with this service.
Root Port	Specifies the port number of the port which offers the lowest cost path from this bridge to the root bridge.
Backbone VPLS	Displays the ID of the B-VPLS

Sample

```
*A:7210-SAS>show>service# id 200 stp

===============================================================================
Stp info, Service 200
===============================================================================
Bridge Id          : 80:00.00:25:ba:08:f6:20  Top. Change Count : 0
Root Bridge        : N/A                      Stp Oper State    : Down
Primary Bridge     : N/A                      Topology Change   : Inactive
Mode               : Rstp                     Last Top. Change  : 0d 00:00:00
Vcp Active Prot.   : N/A
Root Port          : N/A                      External RPC      : 0

===============================================================================
```

```
Stp port info
===============================================================================
Sap/Sdp Id         Oper-     Port-      Port-       Port-  Oper-  Link-  Active
                   State     Role       State       Num    Edge   Type   Prot.
-------------------------------------------------------------------------------
Backbone VPLS      Up        N/A        Forward     2048   N/A    N/A    N/A
1/1/15:200         Up        N/A        Forward     2049   N/A    Pt-pt  N/A
===============================================================================
*A:7210-SAS>show>service#
```

isid-using

Syntax **isid-using** [*ISID*]

Context show>service

Description This command displays services using ISID.

Parameters *ISID* — Displays the service using the specified I-component Service ID (ISID).

Values 0 — 16777215

Label	Description
SvcId	The service identifier.
ISID	Displays the ISID number.
Type	Indicates the type of service.
b-Vpls	Displays the B-VPLS ID.
Adm	Specifies the operating status of the service.
Opr	The current status of the service.
SvcMtu	Indicates the service MTU value.
CustId	Displays the customer ID.

Sample

```
*A:7210-SAS>show>service# isid-using

===============================================================
Services
===============================================================
SvcId      ISID    Type      b-Vpls     Adm  Opr  SvcMtu CustId
---------------------------------------------------------------
100        100     i-VPLS    1          Up   Up   1514   1
```

```
200        200     i-VPLS    2          Up   Up   1514   1
1000       1000    Epipe     2          Up   Up   1514   1
3000       3000    Epipe     1          Up   Up   1514   1
------------------------------------------------------------------
Matching Services : 4
------------------------------------------------------------------
==================================================================
*A:7210-SAS>show>service#
```

i-vpls

Syntax **i-vpls**

Context show>service>id

Description Displays I-VPLS services associated with the B-VPLS service. This command only applies when the service is a B-VPLS.

Output **Show i-vpls Information** — The following table describes service I-vpls output fields.

Label	Description
i-Vpls SvcId	Displays the service ID of the I-VPLS service
Oper ISID	Displays the ISID number.
Admin	Specifies the operating status of the service.
Oper	The current status of the service.

Sample Output

```
*A:7210-SAS>show>service# id 2 i-vpls

===============================================================================
Related i-Vpls services for b-Vpls service 2
===============================================================================
i-Vpls SvcId        Oper ISID           Admin               Oper
-------------------------------------------------------------------------------
200                 200                 Up                  Up
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
===============================================================================
*A:7210-SAS>show>service#
```

epipe

Syntax **epipe**

Context show>service>id

Description This command displays information the Epipe information for the PBB service.

Output **Show Epipe Information** — The following table describes service Epipe output fields.

Label	Description
Epipe SvcId	Displays the service ID of the EPIPE service bound to the B-VPLS service.
Oper ISID	Displays the ISID number.
Admin	Specifies the operating status of the service.
Oper	The current status of the service.

Sample Output

```
*A:7210-SAS>show>service# id 2 epipe

===============================================================================
Related Epipe services for b-Vpls service 2
===============================================================================
Epipe SvcId         Oper ISID           Admin               Oper
-------------------------------------------------------------------------------
1000                1000                Up                  Up
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
===============================================================================
*A:7210-SAS>show>service# id 200 epipe
```

isid-using

Syntax **isid-using [ISID]**

Context show>service

Description This command displays the services using ISID.

Parameters **ISID —** Displays the service using the specified I-component Service ID (ISID).

Values 0 — 16777215

Output **Show Epipe Information** — The following table describes service Epipe output fields.

Label	Description
SvcId	The service identifier.
ISID	Displays the ISID number.
Type	Indicates the type of service.
b-Vpls	Displays the B-VPLS ID.
Admin	Specifies the operating status of the service.
Oper	The current status of the service.
SvcMtu	Indicates the service MTU value.
Customer Id	Displays the customer ID.

Sample Output

```
*A:7210-SAS>show>service# isid-using

===============================================================
Services
===============================================================
SvcId      ISID    Type      b-Vpls     Adm  Opr  SvcMtu CustId
---------------------------------------------------------------
100        100     i-VPLS    1          Up   Up   1514   1
200        200     i-VPLS    2          Up   Up   1514   1
1000       1000    Epipe     2          Up   Up   1514   1
3000       3000    Epipe     1          Up   Up   1514   1
---------------------------------------------------------------
Matching Services : 4
---------------------------------------------------------------
===============================================================
*A:7210-SAS>show>service#
```

service-using

Syntax **service-using [b-vpls] [i-vpls]**

Context show>service

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

Parameters **b-vpls —** Displays matching Epipe services.

i-vpls — Displays matching VPLS instances.

Output **Show Epipe Information** — The following table describes service Epipe output fields.

Label	Description
Service Id	The service identifier.
Type	Indicates the type of service.
Admin	Specifies the operating status of the service.
Oper	The current status of the service.
Customer Id	Displays the customer ID.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.

Sample Output

```
*A:7210-SAS>show>service# service-using b-vpls

===============================================================================
Services [bvpls]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
1            b-VPLS    Up     Up         1                 04/04/2001 23:22:12
2            b-VPLS    Up     Up         1                 04/04/2001 22:13:57
-------------------------------------------------------------------------------
Matching Services : 2
-------------------------------------------------------------------------------
===============================================================================
*A:7210-SAS>show>service#
```

mac-name

Syntax **mac-name** [**detail**]

Context show>service>pbb

Description This command displays information on a specific MAC name.

Parameters **detail** — Displays detail information.

Label	Description
Svc-Id	The service identifier.

ISID	Displays the ISID number.
Name	Displays the MAC name.
Addr	Displays the MAC address

Sample Output

```
*A:7210-SAS>show>service# pbb mac-name test detail

===============================================================================
Services Using MAC name='test' addr='00:25:ba:08:f6:23'
===============================================================================
Svc-Id                              ISID
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:7210-SAS>show>service#
```

PBB 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.

Values service-id: 1 — 214748364

statistics

Syntax **statistics**

Context clear>service>stats

Description This command clears session statistics for this service.

fdb

Syntax **fdb {all | mac** *ieee-address* **| sap** *sap-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.

sap

Syntax **sap** *service-id*

Context clear>service>statistics

Description This command clears statistics for the SAP bound to the service.

Parameters *sap-id* — See Common CLI Command Descriptions on page XXX for command syntax.

counters

Syntax **counters**

Context clear>service>statistics>id

Description This command clears all traffic queue counters associated with the service ID.

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***}**

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.

Parameters **all** — Clears all detected protocol statistics.

sap-id — Clears the specified lease state SAP information.

PBB 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.

Values service-id: 1 — 214748364

event-type

Syntax **[no] event-type {config-change | svc-oper-status-change | sap-oper-status-change}**

Context debug>service>id

Description This command enables a particular debugging event type. The no form of the command disables the event type debugging.

Parameters **config-change —** Debugs configuration change related events.

svc-oper-status-change — Debugs service operational status changes.

sap-oper-status-change — Debugs SAP operational status changes.

sap

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

Context debug>service>id

Description This command enables debugging for a particular SAP.

Parameters *sap-id* — Specifies the SAP ID.

stp

Syntax **stp**

Context debug>service>id

Description This command enables the context for debugging STP.

all-events

Syntax **all-events**

Context debug>service>id>stp

Description This command enables STP debugging for all events.

bpdu

Syntax **[no] bpdu**

Context debug>service>id>stp

Description This command enables STP debugging for received and transmitted BPDUs.

core-connectivity

Syntax **[no] core-connectivity**

Context debug>service>id>stp

Description This command enables STP debugging for core connectivity.

exception

Syntax **[no] exception**

Context debug>service>id>stp

Description This command enables STP debugging for exceptions.

fsm-state-changes

Syntax **[no] fsm-state-changes**

Context debug>service>id>stp

Description This command enables STP debugging for FSM state changes.

fsm-timers

Syntax **[no] fsm-timers**

Context debug>service>id>stp

Description This command enables STP debugging for FSM timer changes.

port-role

Syntax **[no] port-role**

Context debug>service>id>stp

Description his command enables STP debugging for changes in port roles.

port-state

Syntax **[no] port-state**

Context debug>service>id>stp

Description his command enables STP debugging for port states.

sap

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

Context debug>service>id>stp

Description This command enables STP debugging for a specific SAP.

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

Internet Enhanced Service

In This Chapter

This chapter provides information about Internet Enhanced Services when7210 SAS-M and 7210 SAS-T is operated in Network mode and in Access-uplink mode, the process overview, and implementation notes.NOTE: When 7210 SAS-M and 7210 SAS-T is operated in network mode, IES is designed to provide service (or in-band management of the node). When 7210 SAS-M and 7210 SAS-T is operated in access-uplink mode, IES is designed for in-band management of the node.This chapter explicitly notes if a feature is supported in network mode or access-uplink mode.

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.

NOTE: In access-uplink mode, IES is primarily designed for in-band management of the node.

IES allows IP interfaces to participate in the same routing instance used for service network core routing connectivity. IES services require that the IP addressing scheme used by the subscriber be unique between other provider addressing schemes and potentially the entire Internet. While IES is part of the routing domain, the usable IP address space may be limited. This allows a portion of the service provider address space to be reserved for service IP provisioning, and be administered by a separate, but subordinate address authority.

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

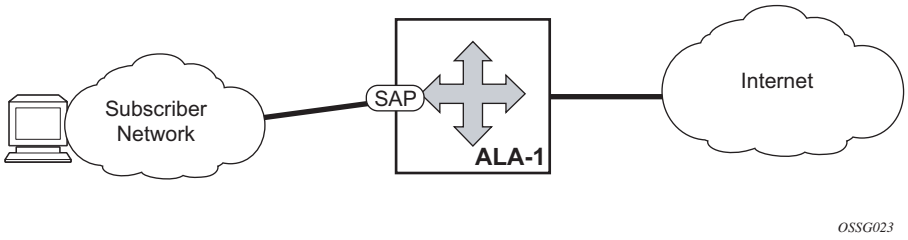

Figure 64: Internet Enhanced Service

The IES service provides in-band management connectivity. Other features include:

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

In access-uplink mode, the IES services provide IP connectivity to the node for in-band management of the node. Most of the management tasks supported with the out-of-band management port are supported with in-band management.

IES Features

This section describes various general 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 options found on the core IP interfaces. The advanced configuration options supported are:

- VRRP - for IES services with more than one IP interface (available only in network mode)
- ICMP Options

In network mode, configuration options found on core IP interfaces not supported on IES IP interfaces are:

- NTP broadcast receipt.

SAPs

Encapsulations

The following SAP encapsulation is supported on IES services in both network mode and access-uplink mode:

- Ethernet null
- Ethernet dot1q
- Ethernet QinQ

In 7210 SAS-M and 7210 SAS-T access-uplink mode, the following access-uplink SAP encapsulations are supported:

- Etherent QinQ (access-uplink QinQ SAP)

Routing Protocols

In network mode, 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:

- OSPF
- Static routing
- IS-IS
- PIM (only on 7210 SAS-M Network mode)

In access-uplink mode, only static routing is supported. Dynamic routing protocols such as OSPF, IS-IS, and others are not supported.

Note that the SAP for the IES IP interface is created at the IES service level, but the routing protocols for the IES IP interface are configured at the routing protocol level for the main router instance.

CPE Connectivity Check

Static routes are used within many IES services. Unlike dynamic routing protocols, there is no way to change the state of routes based on availability information for the associated CPE. CPE connectivity check adds flexibility so that unavailable destinations will be removed from the service provider's routing tables dynamically and minimize wasted bandwidth.

The availability of the far-end static route is monitored through periodic polling. The polling period is configured. If the poll fails a specified number of sequential polls, the static route is marked as inactive.

An ICMP ping mechanism is used to test the connectivity. If the connectivity check fails and the static route is de-activated, the router will continue to send polls and re-activate any routes that are restored.

QoS Policies

When applied to 7210 SAS IES services, service ingress QoS policies only create the unicast meters defined in the policy. The multipoint meters are not created on the service. With IES services, service egress QoS policies function as with other services where the class-based queues are created as defined in the policy.

On 7210 SAS ingress, only meters are supported on all the platforms.

Note: QoS policies only create the unicast meters defined in the policy if PIM is not configured on the associated IP interface; if PIM is configured, the multipoint meters are applied as well.

In access-uplink mode, IES IP interface associated with an access SAP supports use of service ingress QoS policies. IES IP interface associated with an access-uplink SAP does not support use of service ingress QoS policies. IES IP interfaces associated with an access-uplink SAP share the port based ingress and egress QoS policies.

Note that both MAC and IPv4 criteria can be used in the QoS policies for traffic classification in an IES.

CPU QoS for IES interfaces in access-uplink mode

In access-uplink mode, IES IP interface bound to routed VPLS services, IES IP interface on access SAPs and IES IP interface on Access-Uplink SAPs are designed for use with inband management of the node. Consequently, they share a common set of queues for CPU bound management traffic. All CPU bound traffic is policed to pre-defined rates before being queued into CPU queues for

application processing. The system uses meters per application or a set of applications. It does not allocate meters per IP interface. The possibility of CPU overloading has been reduced by use of these mechanisms. Users must use appropriate security policies either on the node or in the network to ensure that this does not happen.

CPU QoS for IES access interfaces in network mode

Traffic bound to CPU received on IES access interfaces are policed/rate-limited and queued into CPU queues. The software allocates a policer per IP application or a set of IP applications, for rate-limiting CPU bound IP traffic from all IES access SAPs. The policers CIR/PIR values are set to appropriate values based on feature scaling and these values are not user configurable. The software allocates a set of queues for CPU bound IP traffic from all IES access SAPs. The queues are either shared by a set of IP applications or in some cases allocated to an IP application. The queues are shaped to appropriate rate based on feature scaling. The shaper rate is not user configurable.

NOTE: The instance of queues and policers used for traffic received on network port IP interfaces is different for traffic received from access port IP interfaces. Additionally the network CPU queues are accorded higher priority than the access CPU queues. This is done to provide better security and mitigate the risk of access traffic affecting network side.

Filter Policies

In network mode, only IP filter policies can be applied to IES services.

In access-uplink mode, only IP filter policies can be applied to IES service when either access SAP or access-uplink SAP is associated with the service.

IPv6 support for IES IP interfaces (applicable for only access-uplink mode)

NOTE: IPv6 addressing is supported for IES IP interfaces in access-uplink mode. IPv6 is not supported with IES IP interfaces in network mode.

In access-uplink mode, IES IP interfaces associated with access-uplink SAPs support IPv6 addressing. IPv6 can be used for in-band management of the node using the IES IP interface.

NOTE: IPv6 IES IP interfaces on access-uplink SAPs is supported only on 7210 SAS-M and 7210 SAS-T in access-uplink mode.

IPv4 and IPv6 route table lookup entries are shared. Before adding routes for IPv6 destinations, route entries in the routed lookup table needs to be allocated for IPv6 addresses. This can be done using the CLI command config> system> resource-profile> max-ipv6-routes. This command allocates route entries for /64 IPv6 prefix route lookups. The system does not allocate any IPv6 route entries by default and user needs to allocate some resources before using IPv6. For the command to take effect the node must be rebooted after making the change. For more information, see the example below and the 7210 SAS Basic System Configuration Guide.

A separate route table is used for IPv6 /128-bit prefix route lookup. A limited amount of IPv6 /128 prefixes route lookup entries is supported. The software enables lookups in this table by default (in other words no user configuration is required to enable Ipv6 /128-bit route lookup).

NOTE: IPv6 interfaces are allowed to be created without allocating IPv6 route entries.

NOTE: IPv6 is not supported for IES IP interfaces associated with access SAPs.

Following features and restrictions is applicable for IPv6 IES IP interfaces:

- IPv6 interfaces supports only static routing.
- Only port-based ingress QoS policies are supported.
- IPv6 filter policies can be used on SAP ingress and egress.
- Routing protocols, such as OSPFv3, and others are not supported.
- A limited amount of IPv6 /128 prefixes route lookup entries is supported on 7210 SAS-M.
- VRRP is not supported.

VRRP support for IES IP interfaces

NOTE: VRRP for IPv4 is supported for IES IP interfaces in network mode only. VRRP is not supported in access-uplink mode. VRRP for IPv6 is not supported.

The Virtual Router Redundancy Protocol (VRRP) for IPv4 is defined in the IETF RFC 3768, Virtual Router Redundancy Protocol. VRRP describes a method of implementing a redundant IP interface shared between two or more routers on a common LAN segment, allowing a group of routers to function as one virtual router. When this IP interface is specified as a default gateway on hosts directly attached to this LAN, the routers sharing the IP interface prevent a single point of failure by limiting access to this gateway address. For more information on use of VRRP, see the "7210 SAS Router Configuration User Guide".

Configuring an IES Service with CLI

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

Topics in this section include:

Basic Configuration

The most basic IES service configuration has the following entities:

- Customer ID (refer to Configuring Customers on page 69)
- An interface to create and maintain IP routing interfaces within IES service ID.
- A SAP on the interface specifying the access port and encapsulation values.

The following example displays a sample configuration of an IES service on ALA-48 on an access-uplink SAP (applicable for access-uplink mode only).

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

The following example displays a basic IES service configuration for Ipv6, along with the use of max-ipv6-routes in 7210 SAS-M and 7210 SAS-T access-uplink mode:

The following displays an example of allocation of IPv6 routes on the node:

```
*A:7210SAS>config>system>res-prof# info
----------------------------------------------
            max-ipv6-routes 1000
----------------------------------------------

NOTE: the node must be rebooted after the above change.

*A:ALA-50>config>service# info
----------------------------------------------
 ies 1000 customer 50 vpn 1000 create
  description "to inband-mgmt"
  interface "to-mgmt" create
   ipv6
    address 10::1/24
     sap 1/1/10:100.* create
    exit
  exit
 no shutdown
----------------------------------------------
*A:ALA-50>config>service#
```

The following example displays a sample configuration of an IES service on ALA-50.

```
*A:ALA-50>config>service# info
----------------------------------------------
```

```
ies 1000 customer 50 vpn 1000 create
     description "to internet"
     interface "to-web" create
          address 10.1.1.1/24
               sap 1/1/10:100 create
          exit
     exit
     no shutdown
----------------------------------------------
*A:ALA-50>config>service#
```

Common Configuration Tasks

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

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

Configuring IES Components

Use the CLI syntax to configure the following entities:

- Configuring an IES Service on page 543
- Configuring IES Interface Parameters on page 544
  - → Configuring SAP Parameters on page 545
  - → Configuring VRRP on page 545

Configuring an IES Service

Use the following CLI syntax to create an IES service:

The following example displays a basic IES service configuration.

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

Configuring IES Interface Parameters

in network mode

```
*A:7210-SAS>config>service>ies>if# info
----------------------------------------------
                arp-timeout 10000
                allow-directed-broadcasts
                icmp
                    ttl-expired 120 38
                exit
                arp-populate
                ip-mtu 1000
                host-connectivity-verify interval 500 timeout 50 retry-count 15
                delayed-enable 150
                bfd 150 receive 300 multiplier 15 echo-receive 3000
                local-proxy-arp
                remote-proxy-arp
                loopback
*A:7210-SAS>config>service>ies>if#
----------------------------------------------
```

The following example displays an IES configuration with interface parameters in access-uplink mode:

```
*A:7210-SAS>config>service>ies>if# info
----------------------------------------------
     arp-timeout 10000
     allow-directed-broadcasts
     icmp
          ttl-expired 120 38
     exit
     ip-mtu 1000
----------------------------------------------
*A:7210-SAS>config>service>ies>if#
```

Configuring SAP Parameters

A SAP is a combination of a port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique within a router.

When configuring IES access SAP parameters, a default QoS policy is applied to each SAP ingress . Additional QoS policies must be configured in the config>qos context. Filter policies are configured in the config>filter context and must be explicitly applied to a SAP. There are no default filter policies.

This example displays an IES SAP configuration.

```
-----------------------------------------------
*A:ALA-A>config>service>ies>if# info
-----------------------------------------------
     address 10.10.36.2/24
     sap 1/1/3:100 create
          ingress
               qos 101
          exit
     exit
-----------------------------------------------
*A:ALA-A>config>service>ies>if#
```

Configuring VRRP

Configuring VRRP parameters on an IES interface is optional and is available only in network mode and is not supported in access-uplink mode. VRRP can be configured in either an owner or non-owner mode. The owner is the VRRP router whose virtual router IP address is the same as the real interface IP address. This is the router that responds to packets addressed to one of the IP addresses for ICMP pings, TCP connections and related addresses. All other virtual router instances participating in this message domain should have the same VRID configured and cannot be configured as an owner.

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 then enter the new information.

The following displays the modified service:

```
*A:ALA-A>config>service>ies# info
----------------------------------------------
     ies 1000 customer 50 create
           description "This is a new description"
           interface "to-web" create
               address 10.1.1.1/24
               mac 00:dc:98:1d:00:00
               sap 1/1/5:0.* create
               exit
           exit
           no shutdown
     exit
----------------------------------------------
*A:ALA-A>config>service#
```

Deleting an IES Service

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

Use the following CLI syntax to delete an IES service:

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

Disabling an IES Service

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

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

Re-Enabling an IES Service

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

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

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

IES Services Command Reference

Command Hierarchies

Global Commands (applicable for both network mode and access-uplink mode)

config
- **service**
  - **ies** *service-id* [**customer** *customer-id*]
  - **no ies** *service-id*
    - **description** *description-string*
    - **no description**
    - **interface**
    - **no interface**
    - **service-name** *service-name*
    - **no service-name**
    - [**no**] **shutdown**

Interface Commands (applicable for network mode)

config
- **service**
  - **ies** *service-id* [**customer** *customer-id*] [**create**]
    - [**no**] **interface** *ip-int-name* [**create**]
      - **address** {*ip-address/mask* | *ip-address netmask*}
      - **no address**
      - **arp-timeout** *seconds*
      - **no arp-timeout**
      - **bfd** *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multi-plier*][**echo-receive** *echo-interval*]
      - **no bfd**
      - **dhcp**

```
                        — description description-string
                        — no description
                        — gi-address ip-address [src-ip-addr]
                        — no gi-address
                        — [no] option
                                — action {replace|drop|keep}
                                — no action
                                — [no] circuit-id [ascii-tuple|ifindex|sap-id|vlan-ascii-
                                  tuple]
                                — [no] remote-id [mac | string string]
                                — [no] vendor-specific-option
                                    — [no] client-mac-address
                                    — [no] sap-id
                                    — [no] service-id
                                    — string text
                                    — no string
                                    — [no] system-id
                        — no server
                        — server server1 [server2...(upto 8 max)]
                        — [no] shutdown
                        — [no] trusted
                — description description-string
                — no description
                — icmp
                        — redirects [number seconds]
                        — no redirects
                        — ttl-expired [number seconds]
                        — no ttl-expired
                        — unreachables [number seconds]
                        — no unreachables
                — ip-mtu octets
                — no ip-mtu
                — [no] loopback
                — [no] sap sap-id [create]
                — [no] shutdown
                — [no] static-arp ip-address [ieee-address]
                — [no] vrrp virtual-router-id
```

Routed VPLS Commands (for devices configured in Access-uplink mode)

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id]
            — interface ip-interface-name [create]
            — no interface-interface-name
                — vpls service-name
                — no vpls
                    — ingress
                        — v4-routed-override-filter ip-filter-id
                        — no v4-routed-override-filter
```

Interface SAP Commands (applicable for network mode)

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id][create]
            — [no] interface ip-int-name
                — [no] sap sap-id [create]
                    — accounting-policy acct-policy-id
                    — no accounting-policy
                    — collect-stats
                    — no collect-stats
                    — description description-string
                    — no description
                    — egress
                        — filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                        — no filter
                    — ingress
                        — [ no ]aggregate-meter-rate rate-in-kbps [burst
                          burst-in-kbits]
                        — filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                        — no filter
                        — meter-override
                        — no meter-override
                            — meter meter-id [create]
                            — no meter meter-id
                                — adaptation-rule [pir adaptation-rule]
                                  [cir adaptation-rule]
                                — cbs size-in-kbytes
                                — no cbs
                                — mbs size-in-kbits
                                — no mbs
                                — mbs mode
                                — no mode
                                — no mode
                                — rate cir cir-rate [pir pir-rate]
                        — [no] qos policy-id
                    — statistics
                        — ingress
                            — counter-mode {in-out-profile-count|forward-
                              drop-count}
                    — [no] tod-suite tod-suite-name
                    — [no] shutdown
```

Interface commands (applicable for access-uplink mode)

```
config
    — service
        — ies service-id [customer customer-id]
            — interface
            — [no] interface ip-int-name
            — address {[ip-address/mask|ip-address netmask] [broadcast all-ones|host-ones]}
            — no address
            — arp-timeout seconds
            — no arp-timeout
            — allow-directed-broadcasts
            — no allow-directed-broadcasts
            — description long description-string
            — no description
            — icmp
                — redirects [number seconds]
                — no redirects
                — ttl-expired [number seconds]
                — no ttl-expired
                — unreachables [number seconds]
                — no unreachables
                — mask-reply
                — no mask-reply
            — ip-mtu octets
            — no ip-mtu
            — [no] ipv6
            — [no] loopback
            — [no]  ieee-address
            — [no] sap sap-id [create]
            — [no] shutdown
            — [no] static-arp ip-address [ieee-address]
```

Interface SAP commands (applicable for access-uplink mode)

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id]
            — [no] interface ip-int-name
                — [no] sap sap-id
                — description description-string
                — no description
            — egress
                — [no] filter [ip ip-filter-id]
                — [no] filter [ ipv6 ip-filter-id]
            — ingress
                — [ no ]aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
                — [no] filter [ip ip-filter-id]
                — [no] filter [ ipv6 ip-filter-id]
                — [no] qos policy-id
            — [no] shutdown
```

VRRP Commands (applicable only for network mode)

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

Interface IPv6 commands (supported only for access-uplink SAPs)

config
- **service**
  - **ies** *service-id* [**customer** *customer-id*] [**create**]
    - [**no**] **interface** *ip-int-name* [**create**]
      - **ipv6**
      - **no ipv6**
        - **[no] address** *ipv6-address/prefix-length [eui-64] [preferred]*
        - **icmp6**
          - **[no] packet-too-big** *number seconds*
          - **[no] param-problem** *number seconds*
          - **[no] redirects** *number seconds*
          - **[no] time-exceeded** *number seconds*
          - **[no] unreachables** *number seconds*
        - **[no] link-local-address** *ipv6-address* **[preferred]**
        - **[no] local-proxy-nd**
        - **[no] neighbor** *ipv6-address mac-address*
        - **[no] proxy-nd-policy** *policy-name* **[policy-name...(upto 5 max)]**

Show Commands

```
show
    — service
        — customer [customer-id] [site customer-site-name]
        — sap-using [sap sap-id]
        — sap-using interface [ip-address | ip-int-name]
        — sap-using [ingress | egress] filter filter-id
        — sap-using [ingress] qos-policy qos-policy-id
        — service-using [ies] [customer customer-id]
        — id service-id
            — all
            — arp [ip-address]|[mac ieee-address]|[sap sap-id]|[interface ip-int-name]
            — base
            — dhcp
                — statistics [sap sap-id] [interface interface-name]
                — summary [interface interface-name | saps]
            — interface [ip-address | ip-int-name] [detail]
```

IES Service Configuration Commands

Generic Commands

shutdown

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

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

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

Special Cases **IES —** The default administrative status of an IES service is down. While the service is down, all its associated virtual router interfaces will be operationally down. The administrative state of the service is not reflected in the administrative state of the virtual router interface.

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

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

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

description

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

Context config>service>ies

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* [**create**]
no ies *service-id*

Context config>service

Description This command creates or edits an IES service instance.

The **ies** command is used to create or maintain an Internet Enhanced Service (IES). If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

IP interfaces defined within the context of an IES service ID must have a SAP created.

When a service is created, the **customer** keyword and *customer-id* must be specified 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 IP interface may be created within a single IES service ID.

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

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

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR, 7450 ESS and 7710 SR on which this service is defined.

Values *service-id*: 1 — 2147483648

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

Values 1 — 2147483647

service-name

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

Context config>service>ies

Description This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a given service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the 7750 SR, 7450 ESS and 7710 SR platforms.

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

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

IES Interface IPv6 Commands

ipv6

Syntax [**no**] **ipv6**

Context config>service>ies>if

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

address

Syntax **address** *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.

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” ICMPv6 messages should be sent. When enabled, ICMPv6 “packet-too-big” messages are generated by this interface.

The **no** form of the command disables the sending of ICMPv6 “packet-too-big” messages.

Default 100 10

Parameters *number* — Specifies the number of “packet-too-big” ICMPv6 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” ICMPv6 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” ICMPv6 messages should be sent. When enabled', “parameter-problem” ICMPv6 messages are generated by this interface.

The **no** form of the command disables the sending of “parameter-problem” ICMPv6 messages.

Default 100 10

number — Specifies the number of “parameter-problem” ICMPv6 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” ICMPv6 messages issued.

Values 1 — 60

Default 10

redirects

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

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

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

When disabled, ICMPv6 redirects are not generated.

Default 100 10

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

Values 10 — 1000

Default 100

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

Values 1 — 60

Default 10

time-exceeded

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

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

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

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

Default 100 10

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

Values 10 — 1000

Default 100

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

Values 1 — 60

Default 10

unreachables

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

link-local-address

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

Context config>service>ies>if>ipv6

Description This command configures the IPv6 link local address.

local-proxy-nd

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

Context config>service>ies>if>ipv6

Description This command enables local proxy neighbor discovery on the interface.

The **no** form of the command disables local proxy neighbor discovery.

proxy-nd-policy

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

Context config>service>ies>if>ipv6

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

Parameters *policy-name* — Specifies an existing neighbor discovery policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains

special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. The specified policy name(s) must already be defined.

neighbor

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

Context config>service>ies>if>ipv6

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

Default none

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

Values x:x:x:x:x:x:x: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 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 Enhanced Service (IES). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.

The **interface** command, under the context of services, is used to create and maintain IP routing interfaces within IES service IDs. The **interface** command can be executed in the context of an IES service ID. The IP interface created is associated with the service core network routing instance and default routing.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config service ies interface** (that is, the network core router instance). Interface names must not be in the dotted decimal notation of an IP address. For example, the name "1.1.1.1" is not allowed, but "int-1.1.1.1" is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All IES IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

The **no** form of this command removes IP the interface and all the associated configuration. The interface must be administratively shutdown before issuing the **no interface** command.

For IES services, the IP interface must be shutdown before the SAP on that interface may be removed.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, $, spaces, 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, an error will occur and context will not be changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*}
address *ip-address mask*
no address

Context config>service>ies>if

Description This command assigns an IP address IP subnet, to an IES IP router interface. Only one IP address can be associated with an IP interface. An IP address must be assigned to each IES IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the 7210 SAS.

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

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

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.

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

/ — The forward slash is a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "/" and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted decimal mask must follow the prefix.

mask-length — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 0 – 30. Note that a mask length of 32 is reserved for system IP addresses.

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask*

parameter indicates the complete mask that will be used in a logical 'AND' function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. Note that a mask of 255.255.255.255 is reserved for system IP addresses.

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>ies>if

Description This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

The **no** form of this command restores **arp-timeout** to the default value.

Default 14400 seconds

Parameters *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

Values

allow-directed-broadcasts

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

Context config>service>ies>if

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

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

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

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

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

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

delayed-enable

Syntax **delayed-enable seconds [init-only]**
no delayed-enable

Context config>service>ies>if

Description This command delays making interface operational by the specified number of seconds. In environments with many subscribers, it can take time to synchronize the subscriber state between peers when the subscriber-interface is enabled (perhaps, after a reboot). To ensure that the state has time to be synchronized, the delayed-enable timer can be specified. The optional parameter init-only can be added to use this timer only after a reboot.

Default no delayed-enable

Parameters *seconds* — Specifies the number of seconds to delay before the interface is operational.

Values 1 — 1200

ip-mtu

Syntax **ip-mtu** *octets*
no ip-mtu

Context config>service>ies>if

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

The MTU that is advertized from the IES size is:

MINIMUM((SdpOperPathMtu - EtherHeaderSize), (Configured ip-mtu))

By default (for Ethernet network interface) if no ip-mtu is configured, the packet size is (1568 - 14) = 1554.

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

Default no ip-mtu

Parameters *octets* — pecifies the number of octets in the IP-MTU.

Values 512 — 9000

loopback

Syntax [**no**] **loopback**

Context config>service>ies>if

Description This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated IES interface cannot be bound to a SAP.

Note that you can configure an IES interface as a loopback interface by issuing the **loopback**

command instead of the **sap** command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

Default none

static-arp

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

Context config>service>ies>if

Description This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP appears in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.

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

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

Default None

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

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

vpls

Syntax vpls *service-name*

Context config>service
config>service>ies>if

Description The vpls command, within the IP interface context, is used to bind the IP interface to the specified service name.

The system does not attempt to resolve the service name provided until the IP interface is placed into the administratively up state (no shutdown). Once the IP interface is administratively up, the system scans the available VPLS services that have the allow-ip-int-binding flag set for a VPLS service associated with the name. If the service name is bound to the service name when the IP interface is already in the administratively up state, the system will immediately attempt to resolve the given name.

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

A VPLS service associated with the specified name that does not have the allow-ip-int-binding flag set or a non-VPLS service associated with the name will be ignored and will not be attached to the IP interface.

If the service name is applied to a VPLS service after the service name is bound to an IP interface and the VPLS service allow-ip-int-binding flag is set at the time the name is applied, the VPLS service is automatically resolved to the IP interface if the interface is administratively up or when the interface is placed in the administratively up state.

If the service name is applied to a VPLS service without the allow-ip-int-binding flag set, the system does not attempt to resolve the applied service name to an existing IP interface bound to the name. To rectify this condition, the flag must first be set and then the IP interface must enter or reenter the administratively up state.

While the specified service name may be assigned to only one service context in the system, it is possible to bind the same service name to more than one IP interface. If two or more IP interfaces are bound to the same service name, the first IP interface to enter the administratively up state (if currently administratively down) or to reenter the administratively up state (if currently administratively up) when a VPLS service is configured with the name and has the allow-ip-intbinding flag set will be attached to the VPLS service. Only one IP interface is allowed to attach to a VPLS service context. No error is generated for the remaining non-attached IP interfaces using the service name.

Once an IP interface is attached to a VPLS service, the name associated with the service cannot be removed or changed until the IP interface name binding is removed. Also, the allow-ip-int-binding flag cannot be removed until the attached IP interface is unbound from the service name. Unbinding the service name from the IP interface causes the IP interface to detach from the VPLS service context. The IP interface may then be bound to another service name or a SAP or SDP binding may be created for the interface using the sap or spoke-sdp commands on the interface.

Default none

Parameters *service-name* — The service-name parameter is required when using the IP interface vpls command and specifies the service name that the system will attempt to resolve to an allow-ip-int-binding enabled VPLS service associated with the name. The specified name is expressed as an ASCII string comprised of up to 32 characters. It does not need to already be associated with a service and the system does not check to ensure that multiple IP interfaces are not bound to the same name.

ingress

Syntax ingress

Context config>service>ies>if>vpls

Description The ingress node in this context under the vpls binding is used to define the routed ip-filter-id optional filter overrides.

v4-routed-override-filter

Syntax v4-routed-override-filter *ip-filter-id*

no v4-routed-override-filter

Context config>service>ies>if>vpls>ingress

Description The v4-routed-override-filter command is used to specify an IP filter ID that is applied to all ingress packets entering the VPLS service. The filter overrides any existing ingress IP filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional and when it is not defined or it is removed, the IP routed packets uses the any existing ingress IP filter on the VPLS virtual port.

The no form of the command is used to remove the IP routed override filter from the ingress IP interface. When removed, the IP ingress routed packets within a VPLS service attached to the IP interface uses the IP ingress filter applied to the packets virtual port when defined.

Default none

Parameters *ip-filter-id* — Specifies the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy in the **configure>filter>ip-filter** context.

Values 1 — 65535

IES Interface ICMP Commands

icmp

Syntax **icmp**

Context config>service>ies>if

Description This command enables the context to configure Internet Control Message Protocol (ICMP) parameters on an IES service

mask-reply

Syntax **[no] mask-reply**

Context config>service>ies>if>icmp

Description 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

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

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

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

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

Default ttl-expired 100 10

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

Values 10 — 1000

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

Values 1 — 60

unreachables

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

Context config>service>ies>if>icmp

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

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

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

sap

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

Context config>service>ies>if

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access uplink port using the **configure port** *port number* **ethernet mode access uplink** command.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service 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 **IES —** A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. Attempts to create a second SAP on an IP interface will fail and generate an error; the original SAP will not be affected.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 987 for command syntax.

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

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

The *port-id* must reference a valid port type. The port must be configured as an uplink access port.

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

IES Filter Commands

filter

Syntax **filter ip** *ip-filter-id*

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

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

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

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

The **no** form of this command removes any configured filter ID association with the SAP. The filter ID itself is not removed from the system.

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.

Values 1 — 65535

egress

Syntax **egress**

Context config>service>ies>if>sap

Description This command enables the context to apply egress policies.

ingress

Syntax **ingress**

Context config>service>ies>if>sap

Description This command enables the context to apply ingress policies

counter-mode

Syntax **counter-mode** {in-out-profile-count|forward-drop-count}

Context config>service>ies>sap>statistics>ingress

Description This command sets the mode of ingress counters associated with the SAP to either octets or packets. On IES SAPs, collect stats cannot be enabled so the mode of the counter can be changed without any reference. Changing the mode of the counter results in loss of previously collected counts and resets the counter. The **no** form of this command is not supported.

Default in-out-profile-count

Parameters **in-out-profile-count —** If the counter mode is specified as "in-out-profile-count", one counter counts the total in-profile packets and octets received on ingress of a SAP and another counts the total out-of-profile packets and octets received on ingress of a SAP. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. Dropped counts are not maintained in hardware when this mode is used. It is obtained by subtracting the sum of in-profile count and out-of-profile count from the total SAP ingress received count and displayed.

forward-drop-count — If the counter mode is specified as "forward-drop-count", one counter counts the forwarded packets and octets received on ingress of a SAP and another counts the dropped packets. The forwarded count is the sum of in-profile and out-of-profile packets/octets received on SAP ingress. The dropped count is count of packets/octets dropped by the policer. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. The in-profile count and out-of-profile count is not individually available when operating in this mode.

tod-suite

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

Context config>service>ies>if>sap

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

Virtual Private Routed Network Service

In This Chapter

This chapter provides information about the Virtual Private Routed Network (VPN) service and implementation notes.VPRN services are supported only in network mode. It is not supported in access-uplink mode.

Topics in this chapter include:

VPRN Service Overview

RFC2547b 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, the label distributed with a VPN route depends on the configured label-mode of the VPRN that is originating the route

Before a customer data packet travels across the service provider's backbone, it is encapsulated with the MPLS label that corresponds, in the customer's VPN, to the route which best matches the packet's destination address. The MPLS packet is further encapsulated with either another MPLS label header, so that it gets tunneled across the backbone to the 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 65 displays a VPRN network diagram example.

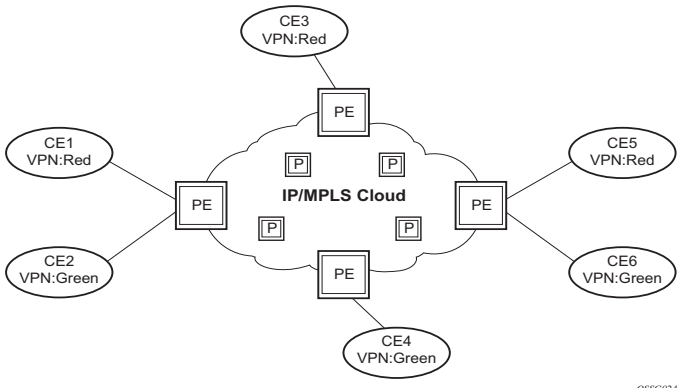

Figure 65: Virtual Private Routed Network

Note: VPRN services is supported only in 7210 SAS devices configured in network mode.

Routing Prerequisites

RFC2547bis requires the following features:

- Multi-protocol extensions
- Extended BGP community support
- BGP capability negotiation
- Parameters defined in RFC 2918

Tunneling protocol options are as follows:

- Label Distribution Protocol (LDP)
- MPLS RSVP-TE tunnels

BGP Support

BGP is used with BGP extensions mentioned in Routing Prerequisites on page 583 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 7210 SAS implementation supports the three (3) type values as defined in the internet draft.

Figure 66: 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 is supported in the service provider core. Multiple sets of route reflectors can be used for different types of BGP routes, including IPv4 and VPN-IPv4. 7210 can only be used a route reflector client. It cannot be used as a route reflector ("server").

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

Each protocol provides controls to limit the number of routes learned from each CE router.

Route Redistribution

Routing information learned from the CE-to-PE routing protocols and configured static routes should be injected in the associated local VPN routing/forwarding (VRF). In the case of dynamic routing protocols, there may be protocol specific route policies that modify or reject certain routes before they are injected into the local VRF.

Route redistribution from the local VRF to CE-to-PE routing protocols is to be controlled via the route policies in each routing protocol instance, in the same manner that is used by the base router instance.

The advertisement or redistribution of routing information from the local VRF to or from the MP-BGP instance is specified per VRF and is controlled by VRF route target associations or by VRF route policies.

VPN-IP routes imported into a VPRN, have the protocol **type bgp-vpn** to denote that it is an VPRN route. This can be used within the route policy match criteria.

CPE Connectivity Check

Static routes are used within many IES and VPRN services. Unlike dynamic routing protocols, there is no way to change the state of routes based on availability information for the associated CPE. CPE connectivity check adds flexibility so that unavailable destinations will be removed from the VPRN routing tables dynamically and minimize wasted bandwidth.

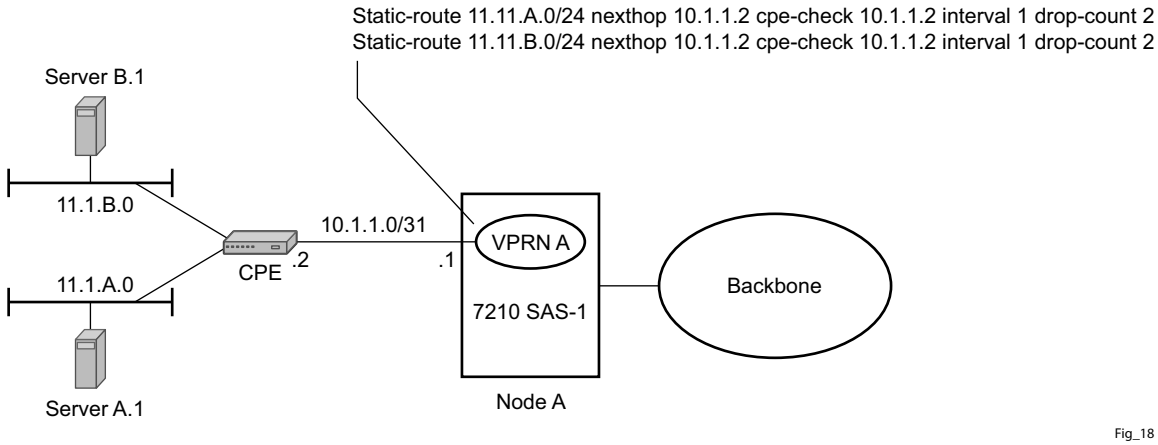

Figure 67: Directly Connected IP Target

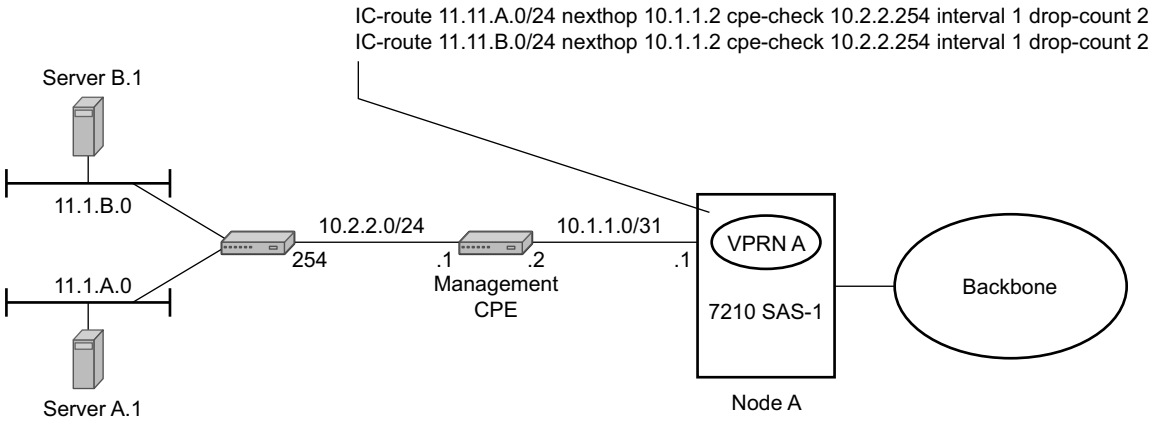

Figure 68: Multiple Hops to IP Target

The availability of the far-end static route is monitored through periodic polling. The polling period is configured. If the poll fails a specified number of sequential polls, the static route is marked as inactive.

Either ICMP ping or unicast ARP mechanism can be used to test the connectivity. ICMP ping is preferred.

If the connectivity check fails and the static route is de-activated, the 7210 SAS router will continue to send polls and re-activate any routes that are restored.

VPRN Features

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

IP Interfaces

VPRN customer IP interfaces can be configured with most of the same options found on the core IP interfaces.

The advanced configuration options supported are:

- VRRP
- ICMP Options

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

- NTP broadcast receipt

SAPs

Encapsulations

The following SAP encapsulations are supported on the 7210 SAS VPRN service:

- Ethernet null
- Ethernet dot1q
- QinQ
- LAG

QoS Policies

When applied to a VPRN SAP, service ingress QoS policies only create the unicast queues defined in the policyif PIM is not configured on the associated IP interface; if PIM is configured, the multipoint queues are applied as well.

Multicast is not supported in VPRN service. Ingress queue, in addition to meters, are supported only on 7210 SAS-X. SAP ingress meters are supported on all other platforms.

For 7210 SAS-M and 7210 SAS-T devices configured in Network mode (with VPRN services), access egress policies are available for use on access ports. Service egress QoS policies are not supported.

Note that both Layer 2 (but dot1p only) or Layer 3 criteria can be used in the QoS policies for traffic classification in an VPRN.

Filter Policies

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

DSCP Marking

DSCP values, dot1p values and forwarding class for all applications is assigned by the system. On ingress, the system uses meters with default values to rate-limit all applications to system defined values. A separate queue and policer is used, one each for all access ports and for all network ports.

Table 24: DSCP/FC Marking

Protocol	IPv4	IPv6	DSCP Marking	Dot1P Marking	Default FC
ARP				7	NC
BGP			48	7	NC
BFD					
PIM (SSM)	Yes				
IGMP	Yes		Yes	Yes	AF
Telnet			34	4	H2
TFTP					
FTP					
SSH (SCP)			34	4	H2
SNMP (get, set, etc.)					
SNMP trap/log					
syslog					
OAM ping					
ICMP ping			0	0	NC
Traceroute			7	0	NC
TACPLUS					
DNS					
SNTP/NTP					
RADIUS					

Default DSCP Mapping Table

```
*A:7210-SAS>show>qos# dscp-table

============================================================
DSCP Mapping
============================================================
DSCP Name      DSCP Value     TOS (bin)      TOS (hex)
------------------------------------------------------------
be             0              0000 0000      00
cp1            1              0000 0100      04
cp2            2              0000 1000      08
cp3            3              0000 1100      0C
cp4            4              0001 0000      10
cp5            5              0001 0100      14
cp6            6              0001 1000      18
cp7            7              0001 1100      1C
cs1            8              0010 0000      20
cp9            9              0010 0100      24
af11           10             0010 1000      28
cp11           11             0010 1100      2C
af12           12             0011 0000      30
cp13           13             0011 0100      34
af13           14             0011 1000      38
cp15           15             0011 1100      3C
cs2            16             0100 0000      40
cp17           17             0100 0100      44
af21           18             0100 1000      48
cp19           19             0100 1100      4C
af22           20             0101 0000      50
cp21           21             0101 0100      54
af23           22             0101 1000      58
cp23           23             0101 1100      5C
cs3            24             0110 0000      60
cp25           25             0110 0100      64
af31           26             0110 1000      68
cp27           27             0110 1100      6C
af32           28             0111 0000      70
cp29           29             0111 0100      74
af33           30             0111 1000      78
cp31           31             0111 1100      7C
cs4            32             1000 0000      80
cp33           33             1000 0100      84
af41           34             1000 1000      88
cp35           35             1000 1100      8C
af42           36             1001 0000      90
cp37           37             1001 0100      94
af43           38             1001 1000      98
cp39           39             1001 1100      9C
cs5            40             1010 0000      A0
cp41           41             1010 0100      A4
cp42           42             1010 1000      A8
cp43           43             1010 1100      AC
cp44           44             1011 0000      B0
cp45           45             1011 0100      B4
ef             46             1011 1000      B8
cp47           47             1011 1100      BC
```

```
nc1             48              1100 0000       C0
cp49            49              1100 0100       C4
cp50            50              1100 1000       C8
cp51            51              1100 1100       CC
cp52            52              1101 0000       D0
cp53            53              1101 0100       D4
cp54            54              1101 1000       D8
cp55            55              1101 1100       DC
nc2             56              1110 0000       E0
cp57            57              1110 0100       E4
cp58            58              1110 1000       E8
cp59            59              1110 1100       EC
cp60            60              1111 0000       F0
cp61            61              1111 0100       F4
cp62            62              1111 1000       F8
cp63            63              1111 1100       FC
===============================================================

default*    0
```

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

CE to PE Routing Protocols

The 7210 SAS VPRN supports the following PE to CE routing protocols:

- BGP
- Static
- OSPF

PE to PE Tunneling Mechanisms

The 7210 SAS supports multiple mechanisms to provide transport tunnels for the forwarding of traffic between PE routers within the 2547bis network.

The 7210 SAS VPRN implementation supports the use of:

- RSVP-TE protocol to create tunnel LSP's between PE routers
- LDP protocol to create tunnel LSP's between PE routers

These transport tunnel mechanisms provide the flexibility of using dynamically created LSPs where the service tunnels are automatically bound (the “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 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 7210 SAS allows setting the maximum number of routes that can be accepted in the VRF for a VPRN service. There are options to specify a percentage threshold at which to generate an event that the VRF table is near full and an option to disable additional route learning when full or only generate an event.

Using OSPF in IP-VPNs

Using OSPF as a CE to PE routing protocol allows OSPF that is currently running as the IGP routing protocol to migrate to an IP-VPN backbone without changing the IGP routing protocol, introducing BGP as the CE-PE or relying on static routes for the distribution of routes into the service providers IP-VPN. The following features are supported:

- Advertisement/redistribution of BGP-VPN routes as summary (type 3) LSAs flooded to CE neighbors of the VPRN OSPF instance. This occurs if the OSPF route type (in the OSPF route type BGP extended community attribute carried with the VPN route) is not external (or NSSA) and the locally configured domain-id matches the domain-id carried in the OSPF domain ID BGP extended community attribute carried with the VPN route.
- OSPF sham links. A sham link is a logical PE-to-PE unnumbered point-to-point interface that essentially rides over the PE-to-PE transport tunnel. A sham link can be associated with any area and can therefore appear as an intra-area link to CE routers attached to different PEs in the VPN.

Service Label Mode of a VPRN

The 7210 SAS allocates one unique (platform-wide) service label per VRF. All VPN-IP routes exported by the PE from a particular VPRN service with that configuration have the same service label. When the PE receives a terminating MPLS packet, the service label value determines the VRF to which the packet belongs. A lookup of the IP packet DA in the forwarding table of the selected VRF determines the next-hop interface.

Configuring a VPRN Service with CLI

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

Topics in this section include:

Basic Configuration

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

- Customer ID (refer to Configuring Customers on page 69)
- 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"
            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
                sap 1/1/10:1 create
                    ingress
                        qos 100
                    exit
                        filter ip 10
                    exit
                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
            no shutdown
----------------------------------------------
*A:ALA-1>config>service>vprn#
```

Common Configuration Tasks

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

1. Associate a VPRN service with a customer ID.
2. Define an autonomous system (optional).
3. Define a route distinguisher (mandatory).
4. Define VRF route-target associations or VRF import/export policies.
5. 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 `config>qos` context).
   → Optional - select filter policies (configured in `config>filter` context).
   → Optional - select accounting policy (configured in `config>log` context).
7. Define BGP parameters (optional).
   → BGP must be enabled in the `config>router>bgp` context.
8. Enable the service.

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`

The following example displays a VPRN service configuration.

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vprn 1 customer 1 create
            route-distinguisher 10001:0
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-1>config>service>vprn#
```

Configuring Global VPRN Parameters

Refer to VPRN Services Command Reference on page 615 for CLI syntax to configure VPRN parameters.

The following example displays a VPRN service with configured parameters.

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vprn 1 customer 1 create
            vrf-import "vrfImpPolCust1"
            vrf-export "vrfExpPolCust1"
            autonomous-system 10000
            route-distinguisher 10001:1
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

Configuring Router Interfaces

Refer to the 7210 SAS OS Router Configuration Guide for command descriptions and syntax information to configure router interfaces.

The following example displays a router interface configurations:

```
ALA48>config>router# info
#------------------------------------------
echo "IP Configuration"
#------------------------------------------
...
        interface "if1"
            address 2.2.2.1/24
        exit
        interface "if2"
            address 10.49.1.46/24
            port 1/1/34
        exit
        interface "if3"
            address 11.11.11.1/24
        exit
...
#------------------------------------------
ALA48>config>router#
```

Configuring VPRN Protocols - BGP

The autonomous system number and router ID configured in the VPRN context only applies to that particular service.

The minimal parameters that should be configured for a VPRN BGP instance are:

- Specify an autonomous system number for the router. See Configuring Global VPRN Parameters on page 603.
- Specify a router ID - Note that if a new or different router ID value is entered in the BGP context, then the new values takes precedence and overwrites the VPRN-level router ID. See Configuring Global VPRN Parameters on page 603.
- 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)
```

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

For more information about the BGP protocol, refer to the 7210 SAS 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.

VPRN BGP CLI Syntax

Use the CLI syntax to configure VPRN BGP parameters (BGP Configuration Commands on page 621).

The following example displays a VPRN BGP configuration:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vprn 1 customer 1 create
            vrf-import "vrfImpPolCust1"
            vrf-export "vrfExpPolCust1"
            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

                        qos 100
                    exit

                        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 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 625 for CLI commands and syntax.

The following example displays a VPRN interface configuration:

```
*A:7210 SAS>config>service>vprn>if# info detail
----------------------------------------------
                no description
                no address
                no mac
                arp-timeout 14400
                no allow-directed-broadcasts
                icmp
                    mask-reply
                    redirects 100 10
                    unreachables 100 10
                    ttl-expired 100 10
                exit
                no arp-populate
                dhcp
                    shutdown
                    no description
                    proxy-server
                        shutdown
                        no emulated-server
                        no lease-time
                    exit
                    no option
                    no server
                    no trusted
                    no lease-populate
                    no gi-address
                    no relay-plain-bootp
                    no use-arp
                exit
                no authentication-policy
                no ip-mtu
                no host-connectivity-verify
                no delayed-enable
                no bfd
                ipcp
                    no peer-ip-address
                    no dns
                exit
                no proxy-arp-policy
                no local-proxy-arp
                no remote-proxy-arp
                no shutdown
----------------------------------------------
*A:7210 SAS>config>service>vprn>if#
```

Configuring a VPRN Interface SAP

A SAP is a combination of a port and encapsulation parameters which identifies the service access point on the interface and within the 7210 SAS. Each SAP must be unique within a router. A SAP cannot be defined if the interface **loopback** command is enabled.

When configuring VPRN interface SAP parameters, a default QoS policy is applied to each ingress and egress SAP. Additional QoS policies and scheduler policies must be configured in the **config>qos** context. Filter policies are configured in the **config>filter** context and must be explicitly applied to a SAP. There are no default filter policies.

The following example displays a VPRN interface SAP configuration:

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vprn 1 customer 1 create
            vrf-import "vrfImpPolCust1"
            vrf-export "vrfExpPolCust1"
            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

                        qos 100
                    exit

                        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#
```

Configuring VPRN Protocols - OSPF

In a VPRN interface, each VPN routing instance is isolated from any other VPN routing instance, and from the routing used across the backbone. OSPF can be run with any VPRN, independently of the routing protocols used in other VPRNs, or in the backbone itself. For more information about the OSPF protocol, see the 7210 SAS OS Routing Protocols Guide.

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

VPRN OSPF CLI Syntax

Refer to Configuring VPRN Protocols - OSPF on page 919 for CLI syntax to configure VPRN parameters.

The following example displays the VPRN OSPF configuration shown above:

```
A:duta>config>service>vprn# info
----------------------------------------------
            router-id 10.10.10.1
            autonomous-system 100
            route-distinguisher 65510:1
            auto-bind ldp
            vrf-target target:65520:1
            interface "to-ixia-1" create
                address 10.1.1.1/24
                sap 1/1/9:1 create
                exit
            exit
            interface "to-ixia-2" create
                address 10.1.2.1/24
                sap 1/1/9:12 create
                exit
            exit
            ospf
                super-backbone
                vpn-domain 0005 0000.0000.0001
                export "from_mbgp_to_ospf"
                area 0.0.0.0
                    interface "to-ixia-2"
                        mtu 1500
                        no shutdown
                    exit
                    sham-link "to-ixia-1" 20.1.1.1
                    exit
                    sham-link "to-ixia-1" 111.11.1.1
                    exit
                exit
            exit
            no shutdown
----------------------------------------------
A:duta>config>service>vprn#
```

For more information about the OSPF protocol, refer to the 7210 SAS OS Routing Protocols Guide.

Service Management Tasks

This section discusses the following service management tasks:

Modifying VPRN Service Parameters

Use the CLI syntax to modify VPRN parameters (VPRN Services Command Reference on page 615).

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"
            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] spoke-sdp sdp-id
               [no] shutdown
```

Disabling a VPRN Service

A VPRN service can be shut down without deleting any service parameters.

CLI Syntax: config>service#
vprn *service-id* [customer *customer-id*]
shutdown

Example: config>service# vprn 1
config>service>vprn# shutdown
config>service>vprn# exit

```
*A:ALA-1>config>service# info
----------------------------------------------
...
        vprn 1 customer 1 create
            shutdown
            vrf-import "vrfImpPolCust1"
            vrf-export "vrfExpPolCust1"
            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

                        qos 100
                    exit
                       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
        exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

Re-enabling a VPRN Service

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

CLI Syntax:

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

VPRN Services Command Reference

Command Hierarchies

VPRN Service Configuration Commands

config
- **service**
  - **vprn** *service-id* [**customer** *customer-id*]
  - **no vprn** *service-id*
    - **auto-bind** {**ldp** | **rsvp-te** | **mpls**}
    - **no auto-bind**
    - **autonomous-system** *as-number*
    - **no autonomous-system**
    - **description** *description-string*
    - **no description**
    - —
    - **maximum-routes** *number* [*log-only*] [**threshold** *percent*]
    - **no maximum-routes**
    - **route-distinguisher** [*ip-address*:*number1* | *asn*:*number2*]
    - **no route-distinguisher**
    - **router-id** *ip-address*
    - **no router-id**
    - [**no**] **shutdown**
    - **snmp-community** *community-name* [**version** *SNMP-version*]
    - **no snmp-community** *community-name*
    - **source-address**
      - **application** *app* [*ip-int-name* | *ip-address*]
      - **no application** *app*
    - [**no**] **spoke-sdp** *sdp-id*
      - **description** *description-string*
      - **no description**
      - [**no**] **shutdown**
    - [**no**] **static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask***}** [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] {**next-hop** *ip-int-name*|*ip-address* | {**cpe-check** *cpe-ip-address* [**interval** *seconds*] [**drop-count** *count*] [**log**]}] {**prefix-list** *prefix-list-name* [**all**|**none**]}]
    - [**no**] **static-route** ***{****ip-prefix/prefix-length* | *ip-prefix netmask***}** [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **indirect** *ip-address* **[cpe-check** *cpe-ip-address* [**interval** *seconds*][**drop-count** *count*] [**log**]] {**prefix-list** *prefix-list-name* [**all**|**none**]}]
    - [**no**] **static-route** ***{****ip-prefix/prefix-length* | *ip-prefix netmask***}** [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **black-hole** {**prefix-list** *prefix-list-name* [**all**|**none**]}]
    - **vrf-export** *policy-name* [*policy-name*...(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**

DHCP Commands

```
config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — dhcp
                — description description-string
                — no description
                — gi-address ip-address [src-ip-addr]
                — no gi-address
                — [no] option
                    — action {replace| drop| keep}
                    — no action
                    — [no] circuit-id [ascii-tuple|ifindex|sap-id|vlan-ascii-tuple]
                    — [no] remote-id [mac | string string]
                    — [no] vendor-specific-option
                        — [no] client-mac-address
                        — [no] sap-id
                        — [no] service-id
                        — string text
                        — no string
                        — [no] system-id
                — no server
                — server server1 [server2...(upto 8 max)]
                — [no] shutdown
                — [no] trusted
```

Interface Commands

config
- **service**
  - **vprn** *service-id* [**customer** *customer-id*]
  - **no vprn** *service-id*
    - [**no**] **interface** *ip-int-name*
      - **address** *ip-address*[/*mask*] [*netmask*] [**broadcast** {**all-ones** | **host-ones**}]
      - **no address**
      - [**no**] **allow-directed-broadcasts**
      - **arp-timeout** [*seconds*]
      - **no arp-timeout**
      - **bfd** *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multi-plier*][**echo-receive** *echo-interval*]
      - **no bfd**
      - **delayed-enable** *seconds*
      - **no delayed-enable**
      - **description** *description-string*
      - **no description** [*description-string*]
      - **icmp**
        - [**no**] **mask-reply**
        - **redirects** *number seconds*
        - **no redirects** [*number seconds*]
        - **ttl-expired** *number seconds*
        - **no ttl-expired** [*number seconds*]
        - **unreachables** *number seconds*
        - **no unreachables** [*number seconds*]
      - [**no**] **local-proxy-arp**
      - [**no**] **loopback**
      - [**no**] **proxy-arp-policy** **policy-name** [*policy-name*...(upto 5 max)]
      - **proxy-arp-policy** *ieee-address*
      - **no proxy-arp-policy**
      - [**no**] **remote-proxy-arp**
      - **static-arp** ieee-address
      - [**no**] **static-arp** [ieee-address]
      - [**no**] **shutdown**
      - **static-arp** ip-address ieee-address
      - [**no**] **static-arp** *ip-address* [*ieee-address*]
      - **[no] vrrp** *virtual-router-id*

Interface VRRP Commands

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

Interface SAP Commands

```
config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — [no] interface ip-int-name [create] [tunnel]
                — [no] sap sap-id
                    — accounting-policy acct-policy-id
                    — no accounting-policy [acct-policy-id]
                    — [no] collect-stats
                    — description description-string
                    — no description [description-string]
                    — Interface Commands
                        — filter ip ip-filter-id
                        — no filter [ip ip-filter-id]
                    — ingress
                        — aggregate-meter-rate rate-in-k bps [burst burst-in-kbits]
                        — no aggregate-meter-rate
                        — filter ip ip-filter-id
                        — no filter [ip ip-filter-id]
                        — meter-override
                        — no meter-override
                            — meter meter-id [create]
                            — no meter meter-id
                                — adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
                                — cbs size-in-kbytes
                                — no cbs
                                — mbs size-in-kbits
                                — no mbs
                                — mbs mode
                                — no mode
                                — no mode
                                — rate cir cir-rate [pir pir-rate]
                        — qos policy-id
                        — no qos [policy-id]
                    — [no] shutdown
                    — statistics
                        — ingress
                            — counter-mode {in-out-profile-count|forward-drop-count}
```

BGP Configuration Commands

```
config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — [no] bgp
                — [no] advertise-inactive
                — [no] aggregator-id-zero
                — always-compare-med {zero | infinity}
                — no always-compare-med
                — [no] as-path-ignore
                — auth-keychain name
                — authentication-key [authentication-key | hash-key] [hash | hash2]
                — no authentication-key
                — best-path-selection
                    — always-compare-med {zero | infinity}
                    — always-compare-med strict-as {zero | infinity}
                    — no always-compare-med
                    — as-path-ignore [ipv4] [vpn-ipv4] [l2-vpn]
                    — no as-path-ignore
                    — ignore-nh-metric
                    — no ignore-nh-metric
                    — ignore-router-id
                    — no ignore-router-id
                — [no] connect-retry seconds
                — [no] damping
                — description description-string
                — no description
                — [no] disable-4byte-asn
                — disable-capability-negotiation
                — no disable-capability-negotiation
                — disable-communities [standard] [extended]
                — no disable-communities
                — [no] disable-fast-external-failover
                — [no] enable-peer-tracking
                — export policy-name [policy-name...(upto 5 max)]
                — no export
                — family [ipv4]
                — no family
                — hold-time seconds [strict]
                — no hold-time
                — import policy-name [policy-name...(up to 5 max)]
                — no import
                — keepalive seconds
                — no keepalive
                — local-preference ip-address
                — no local-preference
                — local-as
                — local-as as-number [private]
                — no local-as
                — local-preference local-preference
                — no local-preference
                — loop-detect {drop-peer | discard-route | ignore-loop| off}
```

- **no loop-detect**
- **med-out** {**number** | **igp-cost**}
- **no med-out**
- **min-as-origination** *seconds*
- **no min-as-origination**
- **min-route-advertisement** *seconds*
- **no min-route-advertisement**
- **multihop** *ttl-value*
- **no multihop**
- **next-hop-self**
- **no next-hop-self**
- **preference** *preference*
- **no preference**
- **peer-as** *as number*
- **no peer-as**
- **[no] path-mtu-discovery**
- [**no**] **rapid-withdrawal**
- [**no**] **remove-private**
- **router-id** *ip-address*
- **no router-id**
- [**no**] **shutdown**
- [**no**] **group** *name* [**dynamic-peer**]
  - [**no**] **advertise-inactive**
  - [**no**] **aggregator-id-zero**
  - [**no**] **as-override**
  - **auth-keychain** *name*
  - **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
  - **no authentication-key**
  - **connect-retry** *seconds*
  - **no connect-retry**
  - [**no**] **damping**
  - **description** *description-string*
  - **no description**
  - [**no**] **disable-4byte-asn**
  - **disable-communities** [**standard**] [**extended**]
  - **no disable-communities**
  - [**no**] **disable-fast-external-failover**
  - [**no**] **enable-peer-tracking**
  - **export** *policy-name* [*policy-name*...(upto 5 max)]
  - **no export**
  - **family** [**ipv4**]
  - **no family**
  - **hold-time** *seconds* [**strict**]
  - **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**
- **peer-as** *as-number*
- **no peer-as**
- **preference** *preference*
- **no preference**
- **[no] path-mtu-discovery**
- **prefix-limit** *limit* [**log-only**] [**threshold** *percent*]
- **no prefix-limit**
- [**no**] **remove-private**
- [**no**] **shutdown**
- **ttl-security** *min-ttl-value*
- **no ttl-security**
- **type** {**internal** | **external**}
- **no type**
- [**no**] **neighbor** *ip-address*
  - [**no**] **advertise-inactive**
  - [**no**] **aggregator-id-zero**
  - [**no**] **as-override**
  - **auth-keychain** *name*
  - **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
  - **no authentication-key**
  - **connect-retry** *seconds*
  - **no connect-retry**
  - [**no**] **damping**
  - **description** *description-string*
  - **no description**
  - [**no**] **disable-4byte-asn**
  - **disable-communities** [**standard**] [**extended**]
  - **no disable-communities**
  - [**no**] **disable-fast-external-failover**
  - [**no**] **enable-peer-tracking**
  - **export** *policy-name* [*policy-name*...(upto 5 max)]
  - **no export**
  - **family** [**ipv4**]
  - **no family**
  - **hold-time** *seconds* [**strict**]
  - **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**
- **peer-as** *as-number*
- **no peer-as**
- **preference** *preference*
- **no preference**
- **[no] path-mtu-discovery**
- **prefix-limit** *limit* [**log-only**] [**threshold** *percent*]
- **no prefix-limit**
- [**no**] **remove-private**
- [**no**] **shutdown**
- **ttl-security** *min-ttl-value*
- **no ttl-security**
- **type** {**internal** | **external**}
- **no type**

OSPF Configuration Commands

```
config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — [no] ospf
                — [no] area area-id
                    — area-range ip-prefix/mask [advertise | not-advertise]
                    — no area-range ip-prefix/mask
                    — [no] blackhole-aggregate
                    — [no] interface ip-int-name [secondary]
                        — [no] advertise-subnet
                        — authentication-key [authentication-key | hash-key] [hash | hash2]
                        — no authentication-key
                        — authentication-type {password | message-digest}
                        — no authentication-type
                        — bfd-enable[remain-down-on-failure]
                        — no bfd-enable
                        — dead-interval seconds
                        — no dead-interval
                        — helper-disable
                        — hello-interval seconds
                        — no hello-interval
                        — interface-type {broadcast | point-to-point}
                        — no interface-type
                        — message-digest-key key-id md5 [key | hash-key] [hash | hash2]
                        — no message-digest-key key-id
                        — metric metric
                        — no metric
                        — mtu bytes
                        — no mtu
                        — [no] passive
                        — priority number
                        — no priority
                        — retransmit-interval seconds
                        — no retransmit-interval
                        — [no] shutdown
                        — transit-delay seconds
                        — no transit-delay
                    — [no] nssa
                        — area-range ip-prefix/mask [advertise | not-advertise]
                        — no area-range ip-prefix/mask
                        — originate-default-route [type-7]
                        — no originate-default-route
                        — [no] redistribute-external
                        — [no] summaries
                    — [no] sham-link ip-int-name ip-address
                        — authentication-key [authentication-key | hash-key] [hash | hash2]
                        — no authentication-key
                        — authentication-type {password | message-digest}
```

- — **no authentication-type**
- — **dead-interval** *seconds*
- — **no dead-interval**
- — **hello-interval** *seconds*
- — **no hello-interval**
- — **message-digest-key** *key-id* **md5** [*key* | *hash-key*] [**hash** | **hash2**]
- — **no message-digest-key** *key-id*
- — **metric** *metric*
- — **no metric**
- — **retransmit-interval** *seconds*
- — **no retransmit-interval**
- — [**no**] **shutdown**
- — **transit-delay** *seconds*
- — **no transit-delay**

— [**no**] **stub**

- — **default-metric** *metric*
- — **no default-metric**
- — [**no**] **summaries**

— [**no**] **virtual-link** *router-id* **transit-area** *area-id*

- — **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
- — **no authentication-key**
- — **authentication-type** {**password** | **message-digest**}
- — **no authentication-type**
- — **dead-interval** *seconds*
- — **no dead-interval**
- — **hello-interval** *seconds*
- — **no hello-interval**
- — **message-digest-key** *key-id* **md5** [*key* | *hash-key*] [**hash** | **hash2**]
- — **no message-digest-key** *key-id*
- — **retransmit-interval** *seconds*
- — **no retransmit-interval**
- — [**no**] **shutdown**
- — **transit-delay** *seconds*
- — **no transit-delay**

— [**no**] **compatible-rfc1583**
— **export** *policy-name* [*policy-name*...(up to 5 max)]
— **no export**
— **external-db-overflow** *limit seconds*
— **no external-db-overflow**
— **external-preference** *preference*
— **no external-preference**
— [**no**] **ignore-dn-bit**
— **import** *policy-name* **[***policy-name***...(upto 5 max)]**
— **no import** *policy-name* **[***policy-name***...(upto 5 max)]**
— **overload** [**timeout** *seconds*]
— **no overload**
— [**no**] **overload-include-stub**
— **overload-on-boot** [**timeout** *seconds*]
— **no overload-on-boot**
— **preference** *preference*
— **no preference**
— **reference-bandwidth** *bandwidth-in-kbps*

- **no reference-bandwidth**
- **router-id** *ip-address*
- **no router-id**
- [**no**] **shutdown**
- [**no**] **super-backbone**
- [**no**] **suppress-dn-bit**
- **timers**
  - [**no**] **lsa-arrival** ***lsa-arrival-time***
  - [**no**] **lsa-generate** *max-lsa-wait* [*lsa-initial-wait* [*lsa-second-wait*]]
  - [**no**] **spf-wait** ***max-spf-wait*** [*spf-initial-wait* [*spf-second-wait*]]
- **vpn-domain** *id* {**0005** | **0105** | **0205** | **8005**}
- **no vpn-domain**
- **vpn-tag** *vpn-tag*
- **no vpn-tag**

Show Commands

show
- **service**
  - **egress-label** *start-label* [*end-label*]
  - **ingress-label** *start-label* [[*end-label*]
  - **id** *service-id*
    - **all**
    - **base**
    - **dhcp**
      - **statistics** [**sap** *sap-id*] [**interface** *interface-name*]
      - **summary** **[interface** *interface-name* | **saps]**
    - **sap** [*sap-id* [**detail**]]
    - **sdp** [*sdp-id* | **far-end** *ip-address*] [**detail**]
  - **labels**
  - **sap-using** [**sap** *sap-id*]
  - **sap-using** **interface** [*ip-address* | *ip-int-name*]
  - **sap-using** [**ingress** | **egress**] **filter** ***f****ilter-id*
  - **sap-using** [**ingress**| **qos-policy** *qos-policy-id*
  - **sdp-using** [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]
  - **sdp-using** [*sdp-id*[:*vc-id*]
  - **service-using** [**vprn**] [**sdp** *sdp-id*] [**customer** *customer-id*]

show

— **router** [*vprn-service-id*]

— **aggregate** [family] [**active**]

— **arp** [<ip-int-name|ip-address[/mask]>|**mac**<ieee-mac-address>|sumary]**[local|dynamic|static|managed]**

— **bgp**

— **auth-keychain [***keychain***]**

— **damping** [*ip-prefix*[/*prefix-length*]] [**decayed|history|suppressed**] [detail] [ipv4]

— **damping** [*ip-prefix*[/*prefix-length*]] [**decayed|history|suppressed**] [detail] vpn-ipv4

— **group** [*name*] [**detail**] inter-as-label

— **neighbor** [*ip-address* [detail]

— **neighbor** [*as-number* [detail]

— **neighbor** [*ip-address* [[**family** *family*] *filter1*][*filter3*]]

— **neighbor** [*as-number* [[**family** *family*] *filter2*]]

— **next-hop [family] [ip-address** [detail]**]**

— **paths**

— **routes** [**family** *family*] [*prefix* [**detail** | **longer**]]

— **routes** [**family** *family*] [*prefix* [**hunt** | **brief**]]

— **routes** [**family** *family*] [**community** *comm-id*]

— **routes** [**family** *family*] [**aspath-regex** *reg-ex*1]

— **routes** [**family**] [*ipv6-prefix*[/*prefix-length*] [**detail** | **longer**]|[**hunt** [**brief**]]]

— **summary** [**all**]

— **interface** [{[*ip-address* | *ip-int-name*] [**detail**]} | **summary** [**family** *family*] [**neighbor** *ip-address*]

— **route-table [family][ip-address[/prefix-length] [longer|exact]]|[protocol protocol-name]|[summary]|**

— **static-arp** [*ip-address* | *ip-int-name* | **mac** *ieee-mac-address*]

— **static-route** [*ip-prefix /mask*] | [**preference** *preference*] | [**next-hop** *ip-address*] [**detail**]

— **tunnel-table** [*ip-address*[/*mask*] [**protocol** *protocol* | **sdp** *sdp-id*]

— **tunnel-table** [**summary**]

Clear Commands

clear
- **router**
  - **bgp**
    - **damping** [{***prefix/mask*** [**neighbor** ***ip-address***]} | {**group** ***name***}]
    - **flap-statistics** [[***ip-prefix/mask***] [**neighbor** ***ip-address***]] | [**group** ***group-name***] | [**regex** ***reg-exp***] | [**policy** ***policy-name***]
    - **neighbor** {***ip-address*** **|** **as** ***as-number*** **| external | all**} [**soft | soft-inbound | statistics**]
    - **protocol**
  - **forwarding-table** [*slot-number*]
  - **interface** [*ip-int-name* | *ip-address*] [**icmp**] [statistics]

clear
- **service**
  - **id** *service-id*
    - **spoke-sdp** *sdp-id***:***vc-id* **ingress-vc-label**
  - **statistics**
    - **sap** *sap-id* {**all** | **counters** | **stp**}
    - **sdp** *sdp-id* **keep-alive**
    - **id** *service-id*
      - **counters**
      - **spoke-sdp** *sdp-id***:***vc-id* {**all** | **counters** | **stp**}
      - **spoke-sdp**

Debug Commands

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

VPRN Service Configuration Commands

Generic Commands

shutdown

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

Context config>service>vprn
config>service>vprn>if
config>service>vprn>if>sap
config>service>vprn>if>sap>static-host
config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
config>service>vprn>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.

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 — This command disables the BGP instance on the given IP interface. Routes learned from a neighbor that is shutdown are immediately removed from the BGP database and RTM. If BGP is globally shutdown, then all group and neighbor interfaces are shutdown operationally. If a BGP group is shutdown, all member neighbor interfaces are shutdown operationally. If a BGP neighbor is shutdown, just that neighbor interface is operationally shutdown.

description

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

Context config>service>vprn>bgp
config>service>vprn
config>service>vprn>if
config>service>vprn>if>sap
config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

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*] [**create**]
no vprn *service-id*

Context config>service

Description This command creates or edits a Virtu9al 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 identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7210 SAS on which this service is defined.

Values *service-id*: 1 — 2147483648
svc-name: 64 characters maximum

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

Values 1 — 2147483647

auto-bind

Syntax **auto-bind** {**ldp | rsvp-te | mpls**}
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.

rsvp-te — Specifies RSVP-TE to be the automatic binding for the SDP assigned to the service

mpls — Specifies that both LDP and RSVP-TE can be used to resolve the BGP nexthop for VPRN routes in an associated VPRN instance.

autonomous-system

Syntax **autonomous-system** *as-number*
no autonomous-system

Context config>service>vprn

Description This command defines the autonomous system (AS) to be used by this VPN routing/forwarding (VRF). This command defines the autonomous system to be used by this VPN routing

The **no** form of the command removes the defined AS from this VPRN context.

Default no autonomous-system

Parameters *as-number* — Specifies the AS number for the VPRN service.

Values 1 — 4294967295

export-limit

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

Context config>service>vprn>grt-lookup

Description This command provides the ability to limit the total number of routes exported from the VRF to the GRT. The value zero (0) provides an override that disables the maximum limit. Setting this value to zero (0) will not limit the number of routes exported from the VRF to the GRT. Configuring a range of one (1) to 1000 will limit the number of routes to the specified value.

The **no** form of the command sets the export-limit to a default of five (5).

Default export-limit 5

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

Values 0 — 1000

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 remote routes that can be held within a VPN routing/ forwarding (VRF) context. Note that **local**, **host**, **static** and **aggregate** routes are not counted.

Note that the VPRN service ID must be in a shutdown state in order to modify maximum-routes command parameters.

If the **log-only** parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then the offending RIP peer (if applicable) is brought down (but the VPRN instance remains up). BGP peering will remain up but the exceeding BGP routes will not be added to the VRF.

The maximum route threshold can dynamically change to increase the number of supported routes even when the maximum has already been reached. Protocols will resubmit their routes which were initially rejected.

The **no** form of the command disables any limit on the number of routes within a VRF context. Issue the **no** form of the command only when the VPRN instance is shutdown.

Default 0 or disabled — The threshold will not be raised.

Parameters *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 set. There are two warnings, the first is a mid-level warning at the threshold value set and the second is a high-level warning at level between the maximum number of routes and the mid-level rate ([mid+max] / 2).

Values 0 — 100

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 — Specifies 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>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.

service-name

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

Context config>service>vprn

Description This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a given service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the 7210 SAS platforms.

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

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

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* — Specifies one or more SNMP community names.

version *SNMP-version* **—** Specifies the SNMP version.

Values **v1**, **v2c**, **both**

source-address

Syntax **source-address**

Context config>service>vprn

Description This command enables 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/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] {**next-hop** *ip-int-name* | *ip-address* | **ipsec-tunnel** *ipsec-tunnel-name*} [**bfd-enable** | {**cpe-check** *cpe-ip-address* [**interval** *seconds*] [**drop-count** *count*] [**log**]}]
[**no**] **static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **indirect** *ip-address* [**cpe-check** *cpe-ip-address* [**interval** *seconds*][**drop-count** *count*] [**log**]]
[**no**] **static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **black-hole**

Context config>service>vprn

Description This command creates static route entries within the associated router instance. When configuring a static route, either **next-hop**, **indirect** or **black-hole** must be configured.

The **no** form of the command deletes the static route entry. If a static route needs to be removed when multiple static routes exist to the same destination, then as many parameters to uniquely identify the static route must be entered.

If a CPE connectivity check target address is already being used as the target address in a different static route, then cpe-check parameters must match. If they do not, the new configuration command will be rejected.

If a static-route command is issued with no cpe-check target but the destination prefix/netmask and next-hop matches a static route that did have an associated cpe-check, the cpe-check test will be removed from the associated static route.

Default No static routes are defined**.**

Parameters *ip-prefix* — The destination address of the aggregate route in dotted decimal notation.

Values

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 — 32

netmask — The subnet mask in dotted decimal notation.

Values 0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)

ip-int-name — The name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed with

ip-address — The IP address of the IP interface. The *ip-addr* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values

ipv4-address	a.b.c.d (host bits must be 0)

enable — Static routes can be administratively enabled or disabled. Use the **enable** parameter to re-enable a disabled static route. In order to enable a static route, it must be uniquely identified by the IP address, mask, and any other parameter that is required to identify the exact static route.

The administrative state is maintained in the configuration file.

Default enable

disable — Static routes can be administratively enabled or disabled. Use the **disable** parameter to disable a static route while maintaining the static route in the configuration. In order to enable a static route, it must be uniquely identified by the IP address, mask, and any other parameter that is required to identify the exact static route.

The administrative state is maintained in the configuration file.

Default enable

interval *seconds* — This optional parameter specifies the interval between ICMP pings to the target IP address.

Values 1 —255 seconds

Default 1 seconds

drop-count *count* — This optional parameter specifies the number of consecutive ping-replies that must be missed to declare the CPE down and to de-active the associated static route.

Values Value range: 1 —255

Default 3

log — This optional parameter enables the ability to log transitions between active and in-active based on the CPE connectivity check. Events should be sent to the system log, syslog and SNMP traps.

next-hop [*ip-address* | *ip-int-name*] **—** Specifies the directly connected next hop IP address used to reach the destination. If the next hop is over an unnumbered interface, the *ip-int-name* of the unnumbered interface (on this node) can be configured.

The **next-hop** keyword and the **indirect** or **black-hole** keywords are mutually exclusive. If an identical command is entered (with the exception of either the **indirect** or **black-hole** parameters), then this static route will be replaced with the newly entered command, and unless specified, the respective defaults for preference and metric will be applied.

The *ip-addr* configured here can be either on the network side or the access side on this node. This address must be associated with a network directly connected to a network configured on this node.

ipsec-tunnel *ipsec-tunnel-name* **—** specifies an IPSec tunnel name up to 32 characters in length.

indirect *ip-address* **—** Specifies that the route is indirect and specifies the next hop IP address used to reach the destination.

The configured *ip-addr* is not directly connected to a network configured on this node. The destination can be reachable via multiple paths. The static route remains valid as long as the address configured as the indirect address remains a valid entry in the routing table. Indirect static routes cannot use an ip-prefix/mask to another indirect static route.

The **indirect** keyword and the **next-hop** or **black-hole** keywords are mutually exclusive. If an identical command is entered (with the exception of either the **next-hop** or **black-hole** parameters), then this static route will be replaced with the newly entered command and unless specified the respective defaults for preference and metric will be applied.

The *ip-addr* configured can be either on the network or the access side and is normally at least one hop away from this node.

black-hole — Specifies a black hole route meaning that if the destination address on a packet matches this static route it will be silently discarded.

The **black-hole** keyword is mutually exclusive with either the **next-hop** or **indirect** keywords. If an identical command is entered, with exception of either the **next-hop** or **indirect** parameters, then the static route is replaced with the new command, and unless specified, the respective defaults for **preference** and **metric** are applied.

preference *preference* — The preference of this static route (as opposed to the routes from different sources such as BGP or OSPF), expressed as a decimal integer. When modifying the **preference** value of an existing static route, unless specified, the metric will not change.

If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of which route to use is determined by the configuration of the ECMP command.

Default 5

Values 1 — 255

metric *metric* — The cost metric for the static route, expressed as a decimal integer. This value is used when importing this static route into other protocols such as OSPF. This value is also used to determine the static route to install in the forwarding table: When modifying the 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.

Values 1..4294967295

bfd-enable — Associates the state of the static route to a BFD session between the local system and the configured nexthop. This keyword cannot be configured if the nexthop is **indirect** or a **blackhole** keywords are specified.

cpe-check *target-ip-address* **—** This parameter specifies the IP address of the target CPE device. ICMP pings will be sent to this target IP address. This parameter must be configured to enable the CPE connectivity feature for the associated static route. The target-ip-address cannot be in the same subnet as the static route subnet itself to avoid possible circular references. This option is mutually exclusive with BFD support on a given static route.

Default no cpe-check enabled

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

Syntax **vrf-target** {**ext-community** | **export** *ext-community* | **import** *ext-community*}
no vrf-target

Context config>service>vprn

Description This command facilitates a simplified method to configure the route target to be added to advertised routes or compared against received routes from other VRFs on the same or remote PE routers (via MP-BGP).

BGP-VPN routes imported with a vrf-target statement will use the BGP preference value of 170 when imported from remote PE routers, or retain the protocol preference value of the exported route when imported from other VRFs in the same router.

Specified **vrf-import** or **vrf-export** policies override the **vrf-target** policy.

The no form of the command removes the vrf-target

Default no vrf-target

Parameters *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

<ext-community> : target:{<ip-addr:comm-val>| <2byte-asnumber:ext-comm-val>|<4byte-asnumber:comm-val>}	
ip-addr	a.b.c.d
comm-val	[0..65535]
2byte-asnumber	[0..65535]
ext-comm-val	[0..4294967295]
4byte-asnumber	[0..4294967295]

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.

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 7210 SAS devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; 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 7210 SAS router. If two *sdp-id* bindings terminate on the same 7210 SAS, an error occurs and the second SDP binding is rejected.

Parameters *sdp-id* — 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

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.

address

Syntax **address** *ip-address/mask* | *ip-address netmask*} [**broadcast** [**all-ones** | **host-ones**]
no address

Context config>service>vprn>if

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 7210 SAS.

The local subnet that the **address** command defines must be part of the services address space within the routing context using the **config router service-prefix** command. The default is to disallow the complete address space to services. 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 /31 subnets).

/ — The forward slash is a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "/" and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted decimal mask must follow the prefix.

mask-length — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 0 – 30. Note that a mask length of 32 is reserved for system IP addresses.

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical 'AND' function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. Note that a mask of 255.255.255.255 is reserved for system IP addresses.

broadcast — The optional **broadcast** parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which 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>if

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

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

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

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

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

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

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

bfd

Syntax **bfd** *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier*] [**echo-receive** *echo-interval*]
no bfd

Context config>service>vprn>if
config>service>ies>if

This command specifies the BFD parameters for the associated IP interface. If no parameters are defined the default value are used.

The multiplier specifies the number of consecutive BFD messages that must be missed from the peer before the BFD session state is changed to down and the upper level protocols (OSPF, IS-IS, BGP) 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 10 — 100000

Default 100

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

Values 10 — 100000

Default 100

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

Values 3— 20

Default 3

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

Values 100 — 100000

Default 100

local-proxy-arp

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

Context config>service>vprn>if

Description This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus will become the forwarding point for all traffic between hosts in that subnet. When local-proxy-arp is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

Default no local-proxy-arp

loopback

Syntax [**no**] **loopback**

Context config>service>vprn>if

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

proxy-arp-policy

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

Context config>service>vprn>if

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.

remote-proxy-arp

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

Context config>service>vprn>if

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

static-arp

Syntax [**no**] **static-arp** *ip-address ieee-mac-address*

Context config>service>vprn>if

Description This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP will appear in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface. If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

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

Default none

Parameters *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.

Interface ICMP Commands

icmp

Syntax **icmp**

Context config>service>vprn>if
config>service>vprn>sub-if>grp-if
config>service>vprn>nw-if

Description This command configures Internet Control Message Protocol (ICMP) parameters on a VPRN service.

mask-reply

Syntax [**no**] **mask-reply**

Context config>service>vprn>if>icmp
config>service>vprn>sub-if>grp-if>icmp
config>service>vprn>nw-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
config>service>vprn>nw-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
config>service>vprn>nw-if>icmp#

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

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

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

Default ttl-expired 100 10

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

Values 10 — 1000

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

Values 1 — 60

unreachables

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

Context config>service>vprn>if>icmp
config>service>vprn>sub-if>grp-if>icmp
config>service>vprn>nw-if>icmp#

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>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 7210 SAS. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP 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.

sap ipsec-*id*.**private** | **public**:*tag* — This parameter associates an IPSec group SAP with this interface. This is the public side for an IPSec tunnel. Tunnels referencing this IPSec group in the private side may be created if their local IP is in the subnet of the interface subnet and the routing context specified matches with the one of the interface.

This context will provide a SAP to the tunnel. The operator may associate an ingress and egress QoS policies as well as filters and virtual scheduling contexts. Internally this creates an Ethernet SAP that will be used to send and receive encrypted traffic to and from the MDA. Multiple tunnels can be associated with this SAP. The “tag” will be a dot1q value. The operator may see it as an identifier. The range is limited to 1 — 4094.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 987 for command syntax.

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

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the slot_number/MDA_number/port_number format. For example 2/3 specifies port 3 on MDA 2 in slot .

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

create — Keyword used to create a SAP instance.

split-horizon-group *group-name* **—** Specifies the name of the split horizon group to which the SAP belongs.

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.

accounting-policy

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

Context config>service>vprn>if>sap

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 — 99

collect-stats

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

Context config>service>vprn>if>sap

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

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

Default no collect-stats

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>vprn>if

Description This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

The **no** form of this command restores **arp-timeout** to the default value.

Default 14400 seconds

Parameters *seconds* — The minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 — 65535

delayed-enable

Syntax **delayed-enable** *seconds* [**init-only**]
no delayed-enable

Context config>service>vprn>if

Description This command delays making interface operational by the specified number of seconds.

In environments with many subscribers, it can take time to synchronize the subscriber state between peers when the subscriber-interface is enabled (perhaps, after a reboot). To ensure that the state has time to be synchronized, the **delayed-enable** timer can be specified. The optional parameter **init-only** can be added to use this timer only after a reboot.

Default no delayed-enable

Parameters *seconds* — Specifies the number of seconds to delay before the interface is operational.

Values 1 — 1200

init-only — Delays the initialization of the subscriber-interface to give the rest of the system time to complete necessary tasks such as allowing routing protocols to converge and/or to allow MCS to sync the subscriber information. The delay only occurs immediately after a reboot.

Interface SAP Filter and QoS Policy Commands

egress

Syntax **egress**

Context config>service>vprn>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

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.

filter

Syntax **filter ip** *ip-filter-id*
filter [mac *mac-filter-id***]**
no filter [*ip ip-filter-id***]**
no filter [mac *mac-filter-id***]**
no filter

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

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.

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

aggregate-meter-rate

Syntax **aggregate-meter-rate** *rate-in-kbps* [**burst** *burst-in-kbits*]
no aggregate-meter-rate

Context config>service>vpls>sap>ingress

Description This command allows the user to configure the SAP aggregate policer. The rate of the SAP aggregate policer must be specified by the user. The user can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the traffic on different FCs and determines the destination of the packet. The packet is either forwarded to an identified profile or dropped.

Note: The sum of CIR of the individual FCs configured under the SAP cannot exceed the PIR rate configured for the SAP. Though the 7210 SAS software does not block this configuration, it is not recommended for use.

The table below provides information about the final disposition of the packet based on the operating rate of the per FC policer and the per SAP aggregate policer:

Per FC meter Operating Rate	Per FC Assigned Color	SAP aggregate meter Operating Rate	SAP aggregate meter color	Final Packet Color
Within CIR	Green	Within PIR	Green	Green or In-profile
Within CIR*	Green	Above PIR	Red	Green or In-profile
Above CIR, Within PIR	Yellow	Within PIR	Green	Yellow or Out-of-Profile
Above CIR, Within PIR	Yellow	Above PIR	Red	Red or Dropped
Above PIR	Red	Within PIR	Green	Red or Dropped

Per FC meter Operating Rate	Per FC Assigned Color	SAP aggregate meter Operating Rate	SAP aggregate meter color	Final Packet Color
Above PIR	Red	Above PIR	Red	Red or Dropped

Table 25: Final Disposition of the packet based on per FC and per SAP policer or meter.

Note*: The row number 2 in the above table is not recommended for use. For more information on this, see the Note in the "**aggregate-meter-rate**" description.

When the SAP aggregate policer is configured, per FC policer can be only configured in "trtcm2" mode (RFC 4115).

Note: The meter modes "srtcm" and "trtcm1" are used in the absence of an aggregate meter.

The SAP ingress meter counters increment the packet or octet counts based on the final disposition of the packet.

If ingress Frame-based accounting is used, the SAP aggregate meter rate accounts for the Ethernet frame overhead. The system accounts for 12 bytes of IFG and 8 bytes of start delimiter.

The **no** form of the command removes the aggregate policer from use.

Default no aggregate-meter-rate

Parameters *rate-in-kbps* — Specifies the rate in kilobits per second.

Values 01 — 20000000 | max

Default max

burst <*burst-in-kilobits*> — Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 —2146959

Default 512

Interface VRRP Commands

vrrp

Syntax **vrrp** *virtual-router-id* [**owner**]
no vrrp *virtual-router-id*

Context config>service>vprn>if

Description This command creates or edits a Virtual Router ID (VRID) on the service IP interface. A VRID is internally represented in conjunction with the IP interface name. This allows the VRID to be used on multiple IP interfaces while representing different virtual router instances.

Two VRRP nodes can be defined on an IP interface. One, both, or none may be defined as owner. The nodal context of vrrp 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
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, this command is not required. If the command is re-executed with a different password key defined, the new key will be used immediately. If a no **authentication-key** command is executed, the password authentication key is restored to the default value. The **authentication-key** command may be executed at any time, 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.

best-path-selection

Syntax **best-path-selection**

Context config>service>bgp

Description This command enables path selection configuration.

always-compare-med

Syntax **always-compare-med** {**zero** | **infinity**}
no always-compare-med strict-as {zero | infinity}
no always-compare-med

Context config>service>bgp>best-path-selection

Description This command configures the comparison of BGP routes based on the MED attribute. The default behavior of 7210 SAS (equivalent to the **no** form of the command) is to only compare two routes on the basis of MED if they have the same neighbor AS (the first non-confed AS in the received AS_PATH attribute). Also by default, a route without a MED attribute is handled the same as though it had a MED attribute with the value 0. The **always-compare-med** command without the **strict-as** keyword allows MED to be compared even if the paths have a different neighbor AS; in this case, if neither **zero** or **infinity** is specified, the **zero** option is inferred, meaning a route without a MED is handled the same as though it had a MED attribute with the value 0. When the **strict-as** keyword is present, MED is only compared between paths from the same neighbor AS, and in this case, **zero** or **infinity** is mandatory and tells BGP how to interpret paths without a MED attribute.

Default **no always-compare-med**

Parameters **zero —** Specifies that for routes learned without a MED attribute that a zero (0) value is used in the MED comparison. The routes with the lowest metric are the most preferred.

infinity — Specifies for routes learned without a MED attribute that a value of infinity (2^32-1) is used in the MED comparison. This in effect makes these routes the least desirable.

strict-as — Specifies BGP paths to be compared even with different neighbor AS.

as-path-ignore

Syntax **as-path-ignore** [**ipv4**] [**vpn-ipv4**] [**l2-vpn**]
no as-path-ignore

Context config>service>bgp>best-path-selection

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**

Parameters **ipv4 —** Specifies that the AS-path length will be ignored for all IPv4 routes.

vpn-ipv4 — Specifies that the length AS-path will be ignored for all IPv4 VPRN routes.

l2-vpn — The AS-path length will be ignored for all L2-VPN NLRIs.

ignore-nh-metric

Syntax **ignore-nh-metric**
no ignore-nh-metric

Context config>service>bgp>best-path-selection

This command instructs BGP to disregard the resolved distance to the BGP next-hop in its decision process for selecting the best route to a destination. When configured in the config>router>bgp>best-path-selection context, this command applies to the comparison of two BGP routes with the same NLRI learned from base router BGP peers. When configured in the config>service>vprn context, this command applies to the comparison of two BGP-VPN routes for the same IP prefix imported into the VPRN from the base router BGP instance. When configured in the config>service>vprn>bgp>best-path-selection context, this command applies to the comparison of two BGP routes for the same IP prefix learned from VPRN BGP peers.

The no form of the command (no ignore-nh-metric) restores the default behavior whereby BGP factors distance to the next-hop into its decision process.

Default **no ignore-nh-metric**

ignore-router-id

Syntax **ignore-router-id**
no ignore-router-id

Context config>service>bgp>best-path-selection

Description When the ignore-router-id command is present and the current best path to a destination was learned from EBGP peer X with BGP identifier x and a new path is received from EBGP peer Y with BGP identifier y the best path remains unchanged if the new path is equivalent to the current best path up to the BGP identifier comparison – even if y is less than x. The no form of the command restores the default behavior of selecting the route with the lowest BGP identifier (y) as best.

Default **no ignore-router-id**

–

backup

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

Context config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp

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

init-delay

Syntax **init-delay** *seconds*
no init-delay

Context config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp

Description This command configures a VRRP initialization delay timer.

Default no init-delay

Parameters *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
config>service>vprn>if>ipv6>vrrp

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

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

Default The physical MAC address associated with the Ethernet interface that the SAP is configured on.

Parameters *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
config>service>vprn>if>ipv6>vrrp

Description This command allows the master instance to dictate the master down timer (non-owner context only).

Default no master-int-inherit

message-interval

Syntax **message-interval** {[*seconds*] [**milliseconds** *milliseconds*]}
no message-interval

Context config>service>vprn>if
config>service>vprn>if>ipv6>vrrp

Description This command sets the advertisement timer and indirectly sets the master down timer on the virtual router instance. The message-interval setting must be the same for all virtual routers participating as a virtual router. Any VRRP advertisement message received with an Advertisement Interval field different than the virtual router instance configured message-interval value will be silently discarded.

The message-interval command is available in both non-owner and owner **vrrp** *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. This parameter is not supported on single-slot chassis.

Values 100 — 900

ping-reply

Syntax [**no**] **ping-reply**

Context config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp

Description 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
config>service>vprn>if>ipv6>vrrp

Description This command associates a VRRP priority control policy with the virtual router instance (non-owner context only).

Parameters *vrrp-policy-id* — Specifies a VRRP priority control policy.

Values 1 — 9999

preempt

Syntax **preempt**
no preempt

Context config>service>vprn>if
config>service>vprn>if>ipv6>vrrp

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
config>service>vprn>if>ipv6>vrrp

Description The priority command provides the ability to configure a specific priority value to the virtual router instance. In conjunction with an optional policy command, the base-priority is used to derive the in-use priority of the virtual router instance.

The priority command is only available in the non-owner **vrrp** *virtual-router-id* nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed. For non-owner virtual router instances, if the priority command is not executed, the base-priority will be set to 100.

The **no** form of this command restores the default value of 100 to base-priority.

Parameters *base-priority* — The base-priority parameter configures the base priority used by the virtual router instance. If a VRRP priority control policy is not also defined, the base-priority will be the in-use priority for the virtual router instance.

Values 1 — 254

Default 100

ssh-reply

Syntax [**no**] **ssh-reply**

Context config>service>vprn>if>vrrp

Description This command enables the non-owner master to reply to SSH Requests directed at the virtual router instance's IP addresses. The SSH request can be received on any routed interface. SSH must not have been disabled at the management security level (either on the parental IP interface or based on the SSH source host address). Proper login and CLI command authentication is still enforced.

When ssh-reply is not enabled, SSH packets to non-owner master virtual IP addresses are silently discarded. Non-owner backup virtual routers never respond to SSH regardless of the ssh-reply configuration.

The ssh-reply command is only available in non-owner **vrrp** *virtual-router-id* nodal context. If the ssh-reply command is not executed, SSH packets to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all SSH packets destined to the non-owner virtual router instance IP addresses.

Default no ssh-reply

standby-forwarding

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

Context config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp

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

The **no** form of the command specifies that a standby router should not forward traffic sent to virtual router's MAC address. However, the standby router should forward traffic sent to the standby router's real MAC address.

Default no standby-forwarding

telnet-reply

Syntax [**no**] **telnet-reply**

Context config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp

Description This command enables the non-owner master to reply to TCP port 23 Telnet Requests directed at the virtual router instance's IP addresses. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Proper login and CLI command authentication is still enforced.

When telnet-reply is not enabled, TCP port 23 Telnet packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to Telnet Requests regardless of the telnet-reply configuration.

The telnet-reply command is only available in non-owner **VRRP** nodal context. If the telnet-reply command is not executed, Telnet packets to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all Telnet packets destined to the non-owner virtual router instance IP addresses.

Default no telnet-reply

traceroute-reply

Syntax [**no**] **traceroute-reply**

Context config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp

Description This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.

When this command is enabled, a non-owner master can reply to traceroute requests directed to the virtual router instance IP addresses.

A non-owner backup virtual router never responds to such traceroute requests regardless of the **traceroute-reply** status.

Default no traceroute-reply

Network Interface Commands

network-interface

Syntax **network-interface** *interface-name* [**create**]
no network-interface *interface-name*

Context config>service>vprn

Description This command configures a network interface.

Counter Mode Commands

statistics

Syntax **statistics**

Context config>service>vprn>if>sap

Description This command enables the context to configure the counters associated with SAP ingress.

ingress

Syntax **ingress**

Context config>service>vprn>if>sap>statistics

Description This command enables the context to configure the ingress SAP statistics counter.

counter-mode

Syntax **counter-mode {in-out-profile-count| forward-drop-count}**

Context config>service>vprn>if>sap>statistics>ingress

Description This command allows the user to set the counter mode for the counters associated with sap ingress meters or policers. A pair of counters is available with each meter. These counters count different events based on the counter mode value.

Note: The counter mode can be changed if an accounting policy is associated with a SAP. If the counter mode is changed the counters associated with the meter are reset and the counts are cleared. If an accounting policy is in use when the counter-mode is changed a new record will be written into the current accounting file.

Execute the following sequence of commands to ensure a new accounting file is generated when the counter-mode is changed:

1. Execute the command **config>service>epipe/vpls>sap> no collect-stats**, to disable writing of accounting records.
2. Change the counter-mode to the desired value, execute the command **config>service>epipe/vpls>sap>counter-mode {in-out-profile-count| forward-drop-count}**.
3. Execute the command **config>service>epipe/vpls>sap> collect-stats**, to enable writing of accounting records.

The **no** form of the command restores the counter mode to the default value.

Default in-out-profile-count

Parameters **in-out-profile-count —** If the counter mode is specified as "in-out-profile-count", one counter counts the total in-profile packets and octets received on ingress of a SAP and another counts the total out-of-profile packets and octets received on ingress of a SAP. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. Dropped counts are not maintained in hardware when this mode is used. It is obtained by subtracting the sum of in-profile count and out-of-profile count from the total SAP ingress received count and displayed.

forward-drop-count — If the counter mode is specified as "forward-drop-count", one counter counts the forwarded packets and octets received on ingress of a SAP and another counts the dropped packets. The forwarded count is the sum of in-profile and out-of-profile packets/octets received on SAP ingress. The dropped count is count of packets/octets dropped by the policer. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. The in-profile count and out-of-profile count is not individually available when operating in this mode.

BGP Commands

bgp

Syntax [**no**] **bgp**

Context config>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 group level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default no aggregator-id-zero — BGP adds the AS number and router ID to the aggregator path attribute.

always-compare-med

Syntax **always-compare-med** {**zero** | **infinity**}
no always-compare-med

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.

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-4byte-asn

Syntax [**no**] **disable-4byte-asn**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command disables the use of 4-byte ASNs. It can be configured at all 3 level of the hierarchy so it can be specified down to the per peer basis.

If this command is enabled 4-byte ASN support should not be negotiated with the associated remote peer(s).

The **no** form of the command resets the behavior to the default which is to enable the use of 4-byte ASN.

disable-capability-negotiation

Syntax [**no**] **disable-capability-negotiation**

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command disables the exchange of capbilities. When command is enabled and after the peering is flapped, any new capabilities are not negotiated and will strictly support IPv4 routing exchanges with that peer.

The **no** form of the command removes this command from the configuration and restores the normal behavior.

Default no disable-capability-negotiation

disable-capability-negotiation

Syntax **[no] disable-capability-negotiation**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command disables the exchange of capbilities. When command is enabled and after the peering is flapped, any new capabilities are not negotiated and will strictly support IPv4 routing exchanges with that peer.

The no form of the command removes this command from the configuration and restores the normal behavior.

Default no disable-capability-negotiation

disable-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

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.

family

Syntax **family** [**ipv4**]
no family

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures the IP family capability.

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

Default no family

Parameters *ipv4* — Provisions IPv4 support.

group

Syntax **group** *name* [**dynamic-peer**]
no group

Context config>service>vprn>bgp

Description This command creates a context to configure a BGP peer group.

The **no** form of the command deletes the specified peer group and all configurations associated with the peer group. The group must be shutdown before it can be deleted.

Default None — No peer groups are defined.

Parameters *name* — The peer group name. Allowed values is a string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

dynamic-peer — This flag designates that the given BGP group will be used by BGP peers created dynamically based on subscriber-hosts pointing to corresponding BGP peering policy. There can be only one BGP group with this flag set in any given VPRN. No bBGP neighbous can be manually configured in a BGP group with this flag set.

Default disabled

neighbor

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

Context config>service>vprn>bgp>group

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.

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

family

Syntax **family** [**ipv4**]
no family

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command specifies the address family or families to be supported over BGP peerings in the base router. This command is additive so issuing the **family** command adds the specified address family to the list.

The **no** form of the command removes the specified address family from the associated BGP peerings. If an address family is not specified, then reset the supported address family back to the default.

Default ipv4

Parameters **ipv4 —** Provisions support for IPv4 routing information.

hold-time

Syntax **hold-time** *seconds* [**strict**]
no hold-time

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

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

strict — When this parameter is specified, the advertised BGP hold-time from the far-end BGP peer must be greater than or equal to the specified value.

import

Syntax **import** *policy* [*policy*...]
no import

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

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 *seconds* parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **keepalive** value is generally one-third of the **hold-time** interval. Even though the OS implementation allows the **keepalive** value and the **hold-time** interval to be independently set, under the following circumstances, the configured **keepalive** value is overridden by the **hold-time** value:

If the specified **keepalive** value is greater than the configured **hold-time**, then the specified value is ignored, and the **keepalive** is set to one third of the current **hold-time** value.

If the specified **hold-time** interval is less than the configured **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, the 7210 SAS uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGP peers. This command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of 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 1— 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.

Default **1** — EBGP peers are directly connected.

64 — IBGP

Parameters *ttl-value* — The TTL value, expressed as a decimal integer.

Values 1 — 255

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.

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

path-mtu-discovery

Syntax **[no] path-mtu-discovery**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables path MTU discovery for the associated TCP connections. In doing so, the MTU for the associated TCP session will be initially set to the egress interface MTU. The DF bit will also be set so that if a router along the path of the TCP connection cannot handle a packet of a particular size without fragmenting, it will send back and ICMP message to set the path MTU for the given session to a lower value that can be forwarded without fragmenting.

The no form of the command disables path MTU discovery.

Default no path-mtu-discovery

prefix-limit

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

Context config>service>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 a certain percentage (default is 90% of this limit), an SNMP trap is sent. When the limit is exceeded, the BGP peering is dropped and disabled.

The **no** form of 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

log-only — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, the BGP peering is not dropped.

percent — The threshold value (as a percentage) that triggers a warning message to be sent. The default value is 90%.

rapid-withdrawal

Syntax [**no**] **rapid-withdrawal**

Context config>service>vprn>bgp

Description This command disables the delay (Minimum Route Advertisement) on sending BGP withdrawals. Normal route withdrawals may be delayed up to the mininum route advertisement to allow for efficient packing of BGP updates.

The **no** form of the command removes this command from the configuration and returns withdrawal processing to the normal behavior.

Default no rapid-withdrawal

remove-private

Syntax [**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, sham-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-prefix/prefix-length* [**advertise** | **not-advertise**]
no area-range *ip-prefix*/*mask*
no area-range *ip-prefix*/*mask*

Context config>service>vprn>ospf>area
ospf>service>vprn>nssa

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 *ipv6-prefix/prefix-length —* The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.

Values ipv6-prefix - x:x:x:x:x:x:x: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]

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.

The **no** form of this command removes this option.

Default blackhole-aggregate

interface

Syntax [**no**] **interface** *ip-int-name* [**secondary**]

Context config>service>vprn>ospf>area

Description This command creates a context to configure an OSPF interface.

By default interfaces are not activated in any interior gateway protocol such as OSPF unless explicitly configured.

The **no** form of the command deletes the OSPF interface configuration for this interface. The **shutdown** command in the **config>router>ospf>interface** context can be used to disable an interface without removing the configuration for the interface.

Default **no interface** — No OSPF interfaces are defined.

Parameters *ip-int-name* — The IP interface name. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service vprn interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

If the IP interface name does not exist or does not have an IP address configured an error message will be returned.

If the IP interface exists in a different area it will be moved to this area.

secondary — Allows multiple secondary adjacencies to be established over a single IP interface.

sham-link

Syntax **sham-link** *ip-int-name ip-address*

Context config>service>vprn>ospf>area

Description This command is similar to a virtual link with the exception that metric must be included in order to distinguish the cost between the MPLS-VPRN link and the backdoor.

Parameters *ip-int-name* — The local interface name used for the sham-link. This is a mandatory parameter and interface names must be unique within the group of defined IP interfaces for **config>router>interface**, **config>service>ies>interface** and **config>service>vprn>interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters, the entire string must be enclosed within double quotes. If the IP interface name does not exist or does not have an IP address configured, an error message will be returned.

ip-address — The IP address of the SHAM-link neighbor in IP address dotted decimal notation. This parameter is the remote peer of the sham link's IP address used to set up the SHAM link. This is a mandatory parameter and must be a valid IP address.

advertise-subnet

Syntax [**no**] **advertise-subnet**

Context config>service>vprn>ospf>area>*if*

Description This command enables advertising point-to-point interfaces as subnet routes (network number and mask). When disabled, point-to-point interfaces are advertised as host routes.

Note that this command is not valid in the OSPF3 context.

The **no** form of the command disables advertising point-to-point interfaces as subnet routes meaning they are advertised as host routes.

Default **advertise-subnet** — Advertises point-to-point interfaces as subnet routes.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>service>vprn>ospf>area>*if*
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link

Description This command configures the password used by the OSPF interface or virtual-link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.

Note that this command is not valid in the OSPF3 context.

All neighboring routers must use the same type of authentication and password for 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>*if*
config>service>vprn>ospf>area>virtual-link

Description This command enables authentication and specifies the type of authentication to be used on the OSPF interface, virtual-link, and sham-link.

Note that this command is not valid in the OSPF3 context.

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

bfd-enable

Syntax **bfd-enable [**remain-down-on-failure**]**
no bfd-enable

Context config>service>vprn>ospf>area>*if*
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link

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.

NOTE: BFD is not supported for IPv6 interfaces.

The no form of this command removes BFD from the associated IGP protocol adjacency.

Default no bfd-enable

Parameters *remain-down-on-failure* — Forces adjacency down on BFD failure.

dead-interval

Syntax **dead-interval** *seconds*
no dead-interval

Context config>service>vprn>ospf>area>*if*
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link

Description This command configures the time, in seconds, that OSPF waits before declaring a neighbor router down. If no hello packets are received from a neighbor for the duration of the dead interval, the router is assumed to be down. The minimum interval must be two times the hello interval.

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

Default **40**

Special Cases **OSPF Interface —** If the **dead-interval** configured applies to an interface, then all nodes on the subnet must have the same dead interval.

Virtual Link — If the **dead-interval** configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same dead interval.

Sham-link — If the **dead-interval** configured applies to a sham-link, then the interval on both endpoints of the sham-link must have the same dead interval **.**

Parameters *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>*if*
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link

Description This command configures the interval between OSPF hellos issued on the interface, virtual link, or sham-link.

The hello interval, in combination with the dead-interval, is used to establish and maintain the adjacency. Use this parameter to edit the frequency that hello packets are sent.

Reducing the interval, in combination with an appropriate reduction in the associated **dead-interval**, allows for faster detection of link and/or router failures at the cost of higher processing costs.

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

Default **hello-interval 10** — A 10-second hello interval.

Special Cases **OSPF Interface —** If the **hello-interval** configured applies to an interface, then all nodes on the subnet must have the same hello interval.

Virtual Link — If the **hello-interval** configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same hello interval.

Sham Link — If the hello-interval configured applies to a sham-link, then the interval on both endpoints of the sham-link must have the same hello interval

Parameters *seconds* — The hello interval in seconds expressed as a decimal integer.

Values 1 — 65535

interface-type

Syntax **interface-type** {**broadcast** | **point-to-point**}
no interface-type

Context config>service>vprn>ospf>area>*if*

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>*if*
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link

Description This command configures a message digest key when MD5 authentication is enabled on the interface, virtual-link or sham-link. Multiple message digest keys can be configured.

Note that this command is not valid in the OSPF3 context.

The **no** form of the command removes the message digest key identified by the *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>if
config>service>vprn>ospf>area>sham-link

Description This command configures an explicit route cost metric for the OSPF interface that overrides the metrics calculated based on the speed of the underlying link.

The **no** form of the command deletes the manually configured interface metric, so the interface uses the computed metric based on the **reference-bandwidth** command setting and the speed of the underlying link.

Default **no metric** — The metric is based on **reference-bandwidth** setting and the link speed.

Parameters *metric* — The metric to be applied to the interface expressed as a decimal integer.

Values 1 — 65535

mtu

Syntax **mtu** *bytes*
no mtu

Context config>service>vprn>ospf>area>if

Description This command configures the OSPF packet size used on this interface. If this parameter is not configured OSPF derives the MTU value from the MTU configured (default or explicitly) in the following contexts:

config>port>ethernet
config>port>sonet-sdh>path
config>port>tdm>t3-e3
config>port>tdm>t1-e1>channel-group

If this parameter is configured, the smaller value between the value configured here and the MTU configured (default or explicitly) in an above-mentioned context is used.

To determine the actual packet size add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF (IP) packet MTU configured in this command.

Use the **no** form of this command to revert to default.

Default **no mtu** — Uses the value derived from the MTU configured in the **config>port** context.

Parameters *bytes* — The MTU to be used by OSPF for this logical interface in bytes.

Values 512 — 9198 (9212-14) (Depends on the physical media)

passive

Syntax [**no**] **passive**

Context config>service>vprn>ospf>area>if

Description This command adds the passive property to the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol.

By default, only interface addresses that are configured for OSPF will be advertised as OSPF interfaces. The **passive** parameter allows an interface to be advertised as an OSPF interface without running the OSPF protocol.

While in passive mode, the interface will ignore ingress OSPF protocol packets and not transmit any OSPF protocol packets.

The **no** form of the command removes the passive property from the OSPF interface.

Default Service interfaces defined in **config>router>service-prefix** are passive.

All other interfaces are not passive.

priority

Syntax **priority** *number*
no priority

Context config>service>vprn>ospf>area>if

Description This command configures the priority of the OSPF interface that is used an election of the designated router on on the subnet.

This parameter is only used if the interface is of type broadcast. The router with the highest priority interface becomes the designated router. A router with priority 0 is not eligible to be Designated Router or Backup Designated Router.

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

Default **priority 1**

Parameters *number* — The interface priority expressed as a decimal integer. A value of 0 indicates the router is not eligible to be the Designated Router of Backup Designated Router on the interface subnet.

Values 0 — 255

retransmit-interval

Syntax **retransmit-interval** *seconds*
no retransmit-interval

Context config>service>vprn>ospf>area>if
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link

Description This command specifies the length of time, in seconds, that OSPF will wait before retransmitting an unacknowledged link state advertisement (LSA) to an OSPF neighbor.

The value should be longer than the expected round trip delay between any two routers on the attached network. Once the retransmit-interval expires and no acknowledgement has been received, the LSA will be retransmitted.

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

Default **retransmit-interval 5**

Parameters *seconds* — The retransmit interval in seconds expressed as a decimal integer.

Values 1 — 3600

transit-delay

Syntax **transit-delay** *seconds*
no transit-delay

Context config>service>vprn>ospf>area>if
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link

Description This command configures the estimated time, in seconds, that it takes to transmit a link state advertisement (LSA) on the interface or virtual link or sham-link.

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

Default **transit-delay 1**

Parameters *seconds* — 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 69) 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).

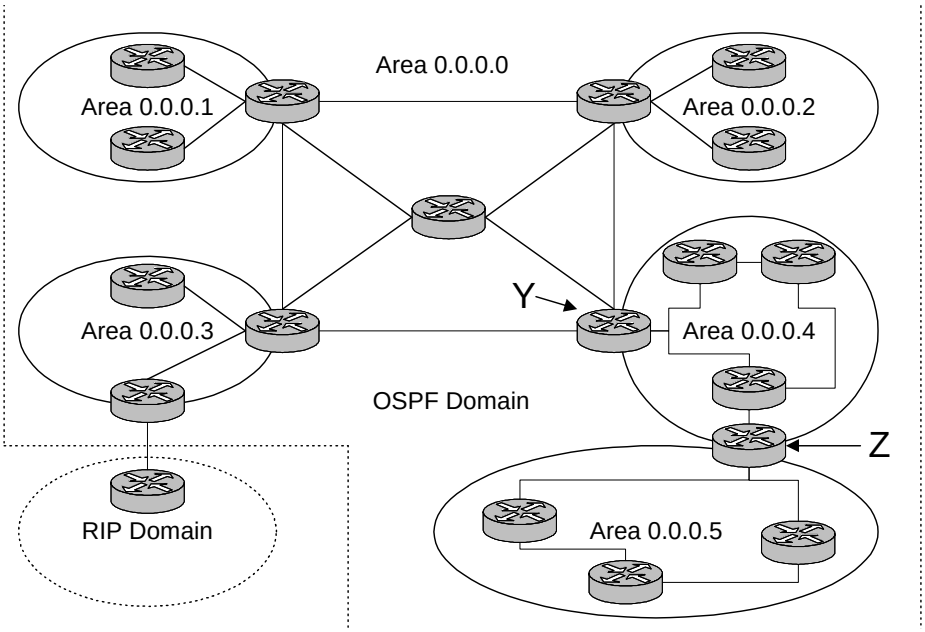

Figure 69: OSPF Areas

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 -1 — 2147483647

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 — 2147483647

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 following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the config>router context.

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

Default **external-preference 150** — OSPF external routes have a default preference of 150.

Parameters *preference* — The preference for external routes expressed as a decimal integer.

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes[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

Route Type	Preference	Configurable (Continued)
BGP	170	Yes

a. Preference for OSPF internal routes is configured with the **preference** command.

Values 1 — 255

ignore-dn-bit

Syntax [**no**] **ignore-dn-bit**

Context config>service>vprn>ospf

Description This command specifies whether to ignore the DN bit for OSPF LSA packets for this instance of OSPF on the router. When enabled, the DN bit for OSPF LSA packets will be ignored. When disabled, the DN bit will not be ignored for OSPF LSA packets.

import

Syntax **import** *policy-name* **[***policy-name***...(upto 5 max)]**
no import

Context config>service>vprn>ospf

Description This command specifies the import route policy to be used to determine which routes are accepted from peers. Route policies are configured in the config>router>policy-options context.

This configuration 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 level is used.

When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be specified. The first policy that matches is applied.

When multiple import commands are issued, the last command entered will override the previous command.

The no form of the command removes the policy association. To remove the association of all policies, use no import without any arguments.

Default no import — No import policy specified.

Parameters *policy-name* — The route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. Route policies are configured in the config>router>policy-options context.

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 the following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the config>router context.

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

Default **preference 10** — OSPF internal routes have a preference of 10.

Parameters *preference* — The preference for internal routes expressed as a decimal integer. Defaults for different route types are listed in the following table.

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes[a]
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

a. Preference for OSPF internal routes is configured with the **preference** command.

Values 1 — 255

reference-bandwidth

Syntax **reference-bandwidth** *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

super-backbone

Syntax [**no**] **super-backbone**

Context config>service>vprn>ospf

Description This command specifies whether CE-PE functionality is required or not. The OSPF super backbone indicates the type of the LSA generated as a result of routes redistributed into OSPF. When enabled, the redistributed routes are injected as summary, external or NSSA LSAs. When disabled, the redistributed routes are injected as either external or NSSA LSAs only.

Default no super-backbone

suppress-dn-bit

Syntax [**no**] **suppress-dn-bit**

Context config>service>vprn>ospf

Description This command specifies whether to suppress the setting of the DN bit for OSPF LSA packets generated by this instance of OSPF on the router. When enabled, the DN bit for OSPF LSA packets generated by this instance of the OSPF router will not be set. When disabled, this instance of the OSPF router will follow the normal proceedure to determine whether to set the DN bit.

Default no suppress-dn-bit

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>timers

Description This command defines the maximum interval between two consecutive SPF calculations in milliseconds. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, then the next SPF will run after 2000 milliseconds, and then next SPF will run after 4000 milliseconds, etc, until it reaches the **spf-wait** value. The SPF interval will stay at the **spf-wait** value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to *spf-initial-wait*.

The timer must be entered in increments of 100 milliseconds. Values entered that do not match this requirement will be rejected.

Use the **no** form of this command to return to the default.

Default no spf-wait

Parameters *max-spf-wait* — Specifies the maximum interval in milliseconds between two consecutive SPF calculations.

Values 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

vpn-domain

Syntax **vpn-domain** [*type* {**0005** | **0105** | **0205** | **8005**}] **id** *id*
no vpn-domain

Context config>service>vprn>ospf

Description This command specifies type of the extended community attribute exchanged using BGP to carry the OSPF VPN domain ID. This applies to VPRN instances of OSPF only. An attempt to modify the value of this object will result in an inconsistent value error when is not a VPRN instance. The parameters are mandatory and can be entered in either order.

Default no vpn-domain

Parameters *id* — Specifies the OSPF VPN domain in the "xxxx.xxxx.xxxx" format. This is exchanged using BGP in the extended community attribute associated with a prefix. This object applies to VPRN instances of OSPF only.

type — Specifies the type of the extended community attribute exchanged using BGP to carry the OSPF VPN domain ID.

Values 0005, 0105, 0205, 8005

vpn-tag

Syntax **vpn-tag** *vpn-tag*
no vpn-tag

Context config>service>vprn>ospf

Description This command specifies the route tag for an OSPF VPN on a PE router. This field is set in the tag field of the OSPF external LSAs generated by the PE. This is mainly used to prevent routing loops. This applies to VPRN instances of OSPF only. An attempt to modify the value of this object will result in an inconsistent value error when is not a VPRN instance.

Default vpn-tag 0

lsa-arrival

Syntax **lsa-arrival** *lsa-arrival-time*
no lsa-arrival

Context config>service>vprn>ospf>timers

Description This parameter defines the minimum delay that must pass between receipt of the same Link State Advertisements (LSAs) arriving from neighbors. It is recommended that the neighbors configured (**lsa-generate**) *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>timers

Description This parameter customizes the throttling of OSPF LSA-generation. Timers that determine when to generate the first, second, and subsequent LSAs can be controlled with this command. Subsequent LSAs are generated at increasing intervals of the *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

Service Global Commands

In This Chapter

This section provides show command descriptions and output.

Show, Clear, Debug, Commands

Services Show Commands

Service Commands

customer

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

Context show>service

Description This command displays service customer information.

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

Default All customer IDs display.

Values 1 — 2147483647

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

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

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

Service-ID	The ID that uniquely identifies a service.
SAP	Specifies the SAP assigned to the service.

Sample Output

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

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

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

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

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

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

Customer-ID : 94043
Contact     : Test Engineer on Duty
Description : TEST Customer
Phone       : (789) 555-1212
------------------------------------------------------
Total Customers : 8
------------------------------------------------------------
*A:ALA-12#
*A:ALA-12# show service customer 274
===============================================================================
```

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

fdb-mac

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

Context show>service

Description This command displays the FDB entry for a given MAC address.

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

expiry — shows amount of time until MAC is aged out.

Sample Output

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


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

```
*A:ALA-48#
```

sdp

Note : SDP commands are not supported by 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

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

Context show>service

Description This command displays SDP information.

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

Parameters *sdp-id* — The SDP ID for which to display information.

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.
Description	Displays a text string describing the SDP.
Admin Path MTU	Displays the desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented. The default value of zero indicates that the path MTU should be computed dynamically from the corresponding MTU of the tunnel.
Opr Path MTU	Displays the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented. In order to be able to bind this SDP to a given service, the value of this object minus the control word size (if applicable) must be equal to or larger than the MTU of the service, as defined by its service MTU.
Far End	Displays the far end IP address.

Label	Description (Continued)
Delivery	The type of delivery used by the SDP: MPLS.
IP address	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Adm Admin State	The desired state of the SDP.
Opr Oper State	The operating state of the SDP.
Flags	Specifies all the conditions that affect the operating status of this SDP.
Signal Signaling	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
Last Status Change	The time of the most recent operating status change to this SDP.
Adv. NTU Over	Specifies whether the advertised MTU of a VLL spoke SDP bind includes the 14-byte L2 header, so that it is backward compatible with pre-2.0 software.
Last Mgmt Change	The time of the most recent management-initiated change to this SDP.
KeepAlive Information	This section displays Keepalive information.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	The length of the SDP echo request messages transmitted on this SDP.
Hello Timeout	The number of seconds to wait for an SDP echo response message before declaring a timeout.
Unmatched Replies	The number of SDP unmatched message replies timer expired.
Max Drop Count	The maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	The amount of time to wait before the keepalive operating status is eligible to enter the alive state.
TX Hello Msgs	The number of SDP echo request messages transmitted since the keepalive was administratively enabled or the counter was cleared.
Rx Hello Msgs	The number of SDP echo request messages received since the keepalive was administratively enabled or the counter was cleared.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field.
Lsp Name	Displays the LSP name.

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

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

Sample Output

```
*A:ALA-7210M# show service sdp
===============================================================================
Services: Service Destination Points
===============================================================================
SdpId    Adm MTU   Opr MTU   IP address       Adm  Opr         Deliver Signal
-------------------------------------------------------------------------------
10       4462      4462      10.20.1.3        Up   Dn NotReady MPLS    TLDP
40       4462      1534      10.20.1.20       Up   Up          MPLS    TLDP
60       4462      1514      10.20.1.21       Up   Up          MPLS    TLDP
100      4462      4462      180.0.0.2        Down Down        MPLS    TLDP
500      4462      4462      10.20.1.50       Up   Dn NotReady MPLS    TLDP
-------------------------------------------------------------------------------
Number of SDPs : 5
===============================================================================
*A:ALA-7210M#


*7210SAS>show>service# sdp 1 detail
```

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

Admin State           : Down                  Oper State         : Down
Signaling             : TLDP                  Metric             : 0
Acct. Pol             : None                  Collect Stats      : Disabled
Last Status Change    : 11/04/2099 22:56:41   Adv. MTU Over.     : No
Last Mgmt Change      : 11/10/2099 15:56:44   VLAN VC Etype      : 0x8100
Bw BookingFactor      : 100                   PBB Etype          : 0x88e7
Oper Max BW(Kbps)     : 0                     Avail BW(Kbps)     : 0
Net-Domain            : default               Egr Interfaces     : n/a
Flags                 : SdpAdminDown NoSysIPAddr
                        TranspTunnDown

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

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

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated

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

sdp-using

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

Context show>service

Description This command displays services using SDP or far-end address options.

Parameters *sdp-id* — Displays only services bound to the specified SDP ID.

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

service-using

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

Context show>service

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

Parameters **epipe —** Displays matching Epipe services.

vpls — Displays matching VPLS instances.

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 command 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
300          Epipe     Up     Up         10                09/05/2006 13:24:15
-------------------------------------------------------------------------------
Matching Services : 2
===============================================================================
*A:ALA-12#
```

eth-ring

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

Context show

Description This command displays the Ethernet rings information.

Parameters *status* — Displays the status information of the Ethernet rings configured on the system.

hierarchy — Displays eth-ring hierarical relationships.

path {a|b} — Displays information related to the configured Ethernet rings.

ring-index — Specifies the ring index of the Ethernet ring.

Values 1—128

Output **Show Ethernet Ring Status —** The following table describes show command output fields.

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

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

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

```
10     Down   Down   a - N/A                -        -        -         -----
                     b - N/A                -        -        -         -----
===============================================================================
Ethernet Tunnel MEP Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM
*A:NS1015C0821>show#
```

Output **Show Ethernet Ring —** The following table describes show command output fields.

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

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

===============================================================================
Ethernet Ring 10 Information
===============================================================================
Description        : (Not Specified)
Admin State        : Down               Oper State       : Down
Node ID            : 00:25:ba:03:48:04
Guard Time         :    5 deciseconds   RPL Node         : rplNone
```

```
Max Revert Time    :  300 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : N/A
Defect Status      :
Sub-Ring Type      : virtualLink        Interconnect-ID  : N/A

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port     Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  -        -                -/-         -               -
  b  -        -                -/-         -               -
===============================================================================
*A:NS1015C0821>show#
```

ETH-CFM Show Commands

eth-cfm

Syntax **eth-cfm**

Context show

Description This command enables the context to display eth-cfm information.

association

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

Context show>eth-cfm

Description This command displays eth-cfm association information.

Parameters *ma-index* — Specifies the maintenance association (MA) index.

Values 1— 4294967295

detail — Displays detailed information for the eth-cfm association.

Output **Show eth-cfm Association Command Output —** The following table describes show eth-cfm association command output fields:

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

Sample Output

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

===============================================================================
CFM Association Table
===============================================================================
Md-index   Ma-index   Name                        CCM-interval Bridge-id
-------------------------------------------------------------------------------
1          1          a1                          1            1
1          2          a2                          1            2
2          1          a1                          1            2
2          2          a2                          1            1
===============================================================================
A:dut-b#
```

cfm-stack-table

Syntax **cfm-stack-table [{all-ports}] [level** <0..7>**] [direction <down>]**
cfm-stack-table port <port-id> **[vlan** <qtag[.qtag]>**] [level** <0..7>**] [direction <down>]**
cfm-stack-table facility [{all-ports|all-lags|all-lag-ports|all-tunnel-meps|all-router-interfaces}] [level <0..7**>] [direction <down>]**
cfm-stack-table facility lag <*id***> [tunnel <**1..4094**>] [level <**0..7**>] [direction <down>]**
cfm-stack-table facility port <*id***> [level <**0..7**>] [direction <down>]**
cfm-stack-table facility router-interface <*ip-int-name***> [level <**0..7**>] [direction <down>]**

Context show>eth-cfm

Description This command displays stack-table information. This stack-table is used to display the various management points MEPs and MIPs that are configured on the system. These can be Service based or facility based. The various option allow the operator to be specific. If no parameters are include then the entire stack-table will be displayed.

Parameters **port** *port-id* **—** Displays the bridge port or aggregated port on which MEPs or MHFs are configured.

vlan *vlan-id* **—** Displays the associated VLAN ID.

level — Display the MD level of the maintenance point.

Values 0 — 7

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

facility — Displays the CFM stack table information for facility MEPs. The base command will display all the facility MEPs. Options may be included in order to further parse the table for specific facility MEP information.

Output **Show eth-cfm CFM Stack Table Command Output —** The following table describes show eth-cfm CFM stack table command output fields:

Label	Description
Sap	Displays associated SAP IDs.
Sdp	Displays the SDP binding for the bridge.
Level Dir	Displays the MD level of the maintenance point.
Md-index	Displays the the maintenance domain (MD) index.
Ma-index	Displays the the maintenance association (MA) index.
Mep-id	Displays the integer that is unique among all the MEPs in the same MA.
Mac-address	Displays the MAC address of the MP.

Sample Output

```
*A:7210SAS>show>eth-cfm# cfm-stack-table

===============================================================================
CFM SAP Stack Table
===============================================================================
Sap            Level Dir  Md-index   Ma-index   Mep-id Mac-address
-------------------------------------------------------------------------------
1/1/18:100     7     Up   7          100        1      00:25:ba:0d:21:13
===============================================================================

===============================================================================
CFM Ethernet Tunnel Stack Table
===============================================================================
Eth-tunnel     Level Dir  Md-index   Ma-index   Mep-id Mac-address
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

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

===============================================================================
CFM Virtual Stack Table
===============================================================================
Service        Level Dir  Md-index   Ma-index   Mep-id Mac-address
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:7210SAS>show>eth-cfm#
```

domain

Syntax **domain** [*md-index*] [**association** *ma-index* | **all-associations**] [**detail**]

Context show>eth-cfm

Description This command displays domain information.

Parameters *md-index* — Displays the index of the MD to which the MP is associated, or 0, if none.

association *ma-index* — Displays the index to which the MP is associated, or 0, if none.

all-associations — Displays all associations to the MD.

detail — Displays detailed domain information.

Output **Show eth-cfm Domain Command Output —** The following table describes show eth-cfm domain command output fields:

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

Sample Output

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

===============================================================================
CFM Domain Table
===============================================================================
Md-index   Level Name                                          Format
-------------------------------------------------------------------------------
1          6     d1                                          charString
2          7     d2                                          charString
===============================================================================
A:dut-b#
```

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**loopback**] [**linktrace**]
mep *mep-id* **domain** *md-index* **association** *ma-index* **remote-mepid** *mep-id* | **all-remote-mepids**
mep *mep-id* **domain** *md-index* **association** *ma-index* **eth-test-results** [**remote-peer** *mac-address*]
mep *mep-id* **domain** *md-index* **association** *ma-index* **one-way-delay-test** [**remote-peer** *mac-address*]
mep*mep-id* **domain** *md-index* **association** *ma-index* **two-way-delay-test** [**remote-peer** *mac-address*]
mep *mep-id* **domain** *md-index* **association** *ma-index* **two-way-slm-test [remote-peer** *macaddress***]**

Context show>eth-cfm

Description This command displays Maintenance Endpoint (MEP) information.

NOTES:

- The show eth-cfm mep mep-id domain md-id association ma-id command does not display CCM ERROR, CCM XCON frames in the output.
- The show eth-cfm mep mep-id domain md-id association ma-id remote-mep rmep-id command does not display some TLVs details.

Parameters *mep-id* — Displays the integer that is unique among all the MEPs in the same MA.

domain *md-index* **—** Displays the index of the MD to which the MP is associated, or 0, if none.

association *ma-index* — Displays the index to which the MP is associated, or 0, if none.

loopback — Displays loopback information for the specified MEP.

linktrace — Displays linktrace information for the specified MEP.

remote-mepid *mep-id* **—** Includes specified remote mep-id information for specified the MEP.

all-remote-mepids — Includes all remote mep-id information for the specified MEP.

eth-test-results — Includes eth-test-result information for the specified MEP.

one-way-delay-test — Includes one-way-delay-test information for the specified MEP.

two-way-delay-test — Includes two-way-delay-test information for the specified MEP.

two-way-slm-test — Includes two-way-slm-test information for the specified MEP.

remote-peer *mac-address* — Includes specified remote mep-id information for the specified MEP.

Sample Output

```
A:dut-b# show eth-cfm mep 1 domain 1 association 1 linktrace
-------------------------------------------------------------------------------
Mep Information
-------------------------------------------------------------------------------
Md-index          : 1                       Direction        : Down
```

```
Ma-index           : 1                        Admin             : Enabled
MepId              : 1                        CCM-Enable        : Enabled
IfIndex            : 35946496                 PrimaryVid        : 1
FngState           : fngReset                 ControlMep        : False
LowestDefectPri    : macRemErrXcon            HighestDefect     : none
Defect Flags       : None
Mac Address        : 00:25:ba:01:c3:6a        CcmLtmPriority    : 7
CcmTx              : 0                        CcmSequenceErr    : 0
Eth-1Dm Threshold  : 3(sec)
Eth-Ais:           : Disabled
Eth-Tst:           : Disabled
CcmLastFailure Frame:
    None
XconCcmFailure Frame:
    None
-------------------------------------------------------------------------------
Mep Linktrace Message Information
-------------------------------------------------------------------------------
LtRxUnexplained    : 0                        LtNextSequence    : 2
LtStatus           : False                    LtResult          : False
TargIsMepId        : False                    TargMepId         : 0
TargMac            : 00:00:00:00:00:00        TTL               : 64
EgressId           : 00:00:00:25:ba:01:c3:6a  SequenceNum       : 1
LtFlags            : useFDBonly
-------------------------------------------------------------------------------
Mep Linktrace Replies
-------------------------------------------------------------------------------
SequenceNum        : 1                        ReceiveOrder      : 1
Ttl                : 63                       Forwarded         : False
LastEgressId       : 00:00:00:25:ba:01:c3:6a  TerminalMep       : True
NextEgressId       : 00:00:00:25:ba:00:5e:bf  Relay             : rlyHit
ChassisIdSubType   : unknown value (0)
ChassisId:
    None
ManAddressDomain:
    None
ManAddress:
    None
IngressMac         : 00:25:ba:00:5e:bf        Ingress Action    : ingOk
IngrPortIdSubType  : unknown value (0)
IngressPortId:
    None
EgressMac          : 00:00:00:00:00:00        Egress Action     : egrNoTlv
EgrPortIdSubType   : unknown value (0)
EgressPortId:
    None
Org Specific TLV:
    None
A:dut-b#
A:dut-b#

A:dut-b# show eth-cfm mep 1 domain 1 association 1 loopback
-------------------------------------------------------------------------------
Mep Information
-------------------------------------------------------------------------------
Md-index           : 1                        Direction         : Down
Ma-index           : 1                        Admin             : Enabled
MepId              : 1                        CCM-Enable        : Enabled
IfIndex            : 35946496                 PrimaryVid        : 1
```

```
FngState           : fngReset                 ControlMep        : False
LowestDefectPri    : macRemErrXcon            HighestDefect     : none
Defect Flags       : None
Mac Address        : 00:25:ba:01:c3:6a        CcmLtmPriority    : 7
CcmTx              : 0                        CcmSequenceErr    : 0
Eth-1Dm Threshold  : 3(sec)
Eth-Ais:           : Disabled
Eth-Tst:           : Disabled
CcmLastFailure Frame:
    None
XconCcmFailure Frame:
    None
-------------------------------------------------------------------------------
Mep Loopback Information
-------------------------------------------------------------------------------
LbRxReply          : 1                        LbRxBadOrder      : 0
LbRxBadMsdu        : 0                        LbTxReply         : 0
LbSequence         : 2                        LbNextSequence    : 2
LbStatus           : False                    LbResultOk        : True
DestIsMepId        : False                    DestMepId         : 0
DestMac            : 00:00:00:00:00:00        SendCount         : 0
VlanDropEnable     : True                     VlanPriority      : 7
Data TLV:
    None
A:dut-b#


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

===============================================================
Eth CFM Two-way Delay Test Result Table
===============================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
---------------------------------------------------------------
00:25:ba:00:5e:bf     507                 507
===============================================================
*A:dut-b#

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

===============================================================
Eth CFM Two-way Delay Test Result Table
===============================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
---------------------------------------------------------------
00:25:ba:00:5e:bf     507                 507
===============================================================
*A:dut-b#

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

===============================================================
Eth CFM ETH-Test Result Table
===============================================================
                                  Current        Accumulate
                  FrameCount      ErrBits        ErrBits
Peer Mac Addr     ByteCount       CrcErrs        CrcErrs
```

```
-------------------------------------------------------------
00:25:ba:01:c3:6a 6              0              0
                  384            0              0
=============================================================
*A:dut-a#

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

=============================================================
Eth CFM ETH-Test Result Table
=============================================================
                                 Current        Accumulate
                  FrameCount     ErrBits        ErrBits
Peer Mac Addr     ByteCount      CrcErrs        CrcErrs
-------------------------------------------------------------
00:25:ba:01:c3:6a 6              0              0
                  384            0              0
=============================================================

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

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


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

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

Show output for two-way-slm-test

```
*A:7210SAS# show eth-cfm mep 1 domain 7 association 100 two-way-slm-test

===============================================================================

Eth CFM Two-way SLM Test Result Table (Test-id: 1)

===============================================================================

Peer Mac Addr      Remote MEP       Count     In Loss    Out Loss         Unack

-------------------------------------------------------------------------------

00:25:ba:0d:1e:12           2           1           0           0            0
```

```
===============================================================================

*A:7210SAS#
```

connection-profile

Syntax **connection-profile [conn-prof-id] [associations]**

Context show

Description This command displays connection profile information.

Parameters *conn-prof-id* — Specifies the connection profile ID.

Values 1 — 8000

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

Output **The following table describes show connection-profile command output fields**

Label	Description
CP Index	Identifies the connection-profile.
Number of Members	Indicates the number of ATM connection profile members not applicable for 7210.
HasRange	Indicates whether VLAN range is configured or not

Sample Output

Show output for connection-profile

```
*7210SAS>show# connection-profile

===============================================================================
Connection Profile Summary Information
===============================================================================
CP Index  Number of HasRange
          Members
-------------------------------------------------------------------------------
1         0         Yes
2         0         Yes
3         0         Yes
5         0         Yes
6         0         Yes
100       0         Yes
200       0         Yes
300       0         Yes
400       0         Yes
500       0         Yes
600       0         Yes
700       0         Yes
800       0         Yes
900       0         Yes
===============================================================================
*7210SAS>show#
```

Show output for connection-profile associations

```
*A:7210SAS>show# connection-profile associations

===============================================================================
Connection Profile Summary Information
===============================================================================
CP Index  Number of HasRange
          Members
-------------------------------------------------------------------------------
1         0         No
===============================================================================
*A:7210SAS>show#
```

IES Show Commands

customer

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

Context show>service

Description This command displays service customer information.

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

Default All customer IDs display

Values 1 — 2147483647

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

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

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

Sample Output

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

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

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

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

Customer-ID : 7
Contact     : Lucy
Description : ABC Customer
Phone       : (567) 555-1212

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

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

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

-------------------------------------------------------
Total Customers : 8
-------------------------------------------------------
*A:ALA-12#


*A:ALA-12# show service customer 274
===============================================================================
Customer  274
===============================================================================
Customer-ID : 274
Contact     : Mssrs. Beaucoup
Description : ABC Company
Phone       : 650 123-4567
```

```
-------------------------------------------------------------------------------
Multi Service Site
-------------------------------------------------------------------------------
Site        : west
Description : (Not Specified)
===============================================================================
*A:ALA-12#


*A:ALA-12# show service customer 274 site west
===============================================================================
Customer  274
===============================================================================
Customer-ID : 274
Contact     : Mssrs. Beaucoup
Description : ABC Company
Phone       : 650 123-4567
-------------------------------------------------------------------------------
Multi Service Site
-------------------------------------------------------------------------------
Site        : west
Description : (Not Specified)
Assignment  : Card 5
I. Sched Pol: SLA1
E. Sched Pol: (Not Specified)
-------------------------------------------------------------------------------
Service Association
-------------------------------------------------------------------------------
No Service Association Found.
===============================================================================
*A:ALA-12#
```

sap-using

Syntax **sap-using** [**sap** *sap-id*]
sap-using interface [*ip-address* | *ip-int-name*]
sap-using [**ingress** | **egress**] **filter** *filter-id*
sap-using [**ingress**] **qos-policy** *qos-policy-id*

Context show>service

Description Displays SAP information.

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

Parameters **sap *sap-id* —** Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 987 for command syntax.

ingress — Specifies matching an ingress policy.

egress — Specifies matching an egress policy.

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

Values 1 — 65535

interface — Specifies matching SAPs with the specified IP interface.

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

Values 1.0.0.0 — 223.255.255.255

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

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

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

Sample Output

```
*A:DUT-B# show service sap-using sap 1/1/3:100.*
===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.   Adm  Opr
                                           QoS   Fltr    Fltr
-------------------------------------------------------------------------------
1/1/1                           6          1     none    none   Up   Down
1/1/2                           700        1     none    none   Up   Down
-------------------------------------------------------------------------------
Number of SAPs : 2
===============================================================================
*A:DUT-B#
```

service-using

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

Context show>service

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

Parameters **ies —** Displays matching IES services.

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

Default Services associated with an customer.

Values 1 — 2147483647

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

Label	Description
Service Id	The value that identifies the service.
Type	Specifies the service type configured for the service ID.
Adm	The administrative state of the service.
Opr	The operating state of the service.
CustomerID	The ID of the customer who owns this service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this service.

Sample Output

```
A:ALA-48# show service service-using ies
===============================================================================
Services [ies]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
88           IES       Up     Down       8                 07/25/2006 15:46:28
89           IES       Up     Down       8                 07/25/2006 15:46:28
104          IES       Up     Down       1                 07/25/2006 15:46:28
200          IES       Up     Down       1                 07/25/2006 15:46:28
214          IES       Up     Down       1                 07/25/2006 15:46:28
321          IES       Up     Down       1                 07/25/2006 15:46:28
322          IES       Down   Down       1                 07/25/2006 15:46:28
1001         IES       Up     Down       1730              07/25/2006 15:46:28
-------------------------------------------------------------------------------
Matching Services : 8
-------------------------------------------------------------------------------
A:ALA-48#
```

id

Syntax **id** *service-id* {**all** | **arp** | **base** | **sap**| **interface | mstp-configuration** }

Context show>service

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

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

all — Display detailed information about the service.

arp — Display ARP entries for the service.

base — Display basic service information.

interface — Display service interfaces.

mstp-confi — guration - Display MSTP information.

sap — Display SAPs associated to the service.

split-horizon-group — Display split horizon group information.

all

Syntax **all**

Context show>service>id

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

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

Label	Description
Service Detailed Information	
`Service Id`	The service identifier.
`VPN Id`	The number which identifies the VPN.
`Service Type`	Specifies the type of service.
`SDP Id`	The SDP identifier.
`Description`	Generic information about the service.
`Customer Id`	The customer identifier.
`Last Mgmt Change`	The date and time of the most recent management-initiated change to this customer.
`SAP Count`	The number of SAPs specified for this service.
`SDP Bind Count`	The number of SDPs bound to this service.

Label	Description (Continued)
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.
Type	Indicates whether this Service SDP binding is a spoke or a mesh.
Admin Path MTU	The largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the service.
Oper State	The current status of the service.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.

Label	Description (Continued)
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far-end field.
Number of SDPs	The total number SDPs applied to this service ID.
Service Access Points	
Service Id	The service identifier.
Port Id	The ID of the access port where this SAP is defined.
Description	Generic information about the SAP.
Encap	The value of the label used to identify this SAP on the access port.
Admin State	The desired state of the SAP.
Oper State	The operating state of the SAP.
Last Changed	The date and time of the last change.
Admin MTU	The largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The SAP ingress QoS policy ID.
Egress qos-policy	The SAP egress QoS policy ID.
Ingress Filter-Id	The SAP ingress filter policy ID.
Egress Filter-Id	The SAP egress filter policy ID.
Multi Svc Site	Indicates the multi-service site that the SAP is a member.
Ingress sched-policy	Indicates the ingress QoS scheduler for the SAP.
Egress sched-policy	Indicates the egress QoS scheduler for the SAP.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
SAP Statistics	
Dropped	The number of packets or octets dropped.
Offered Hi Priority	The number of high priority packets, as determined by the SAP ingress QoS policy.

Label	Description (Continued)
Offered Low Priority	The number of low priority packets, as determined by the SAP ingress QoS policy.
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Split Horizon Group Specifics	
Split Horizon Group	Displays the name of the split horizon group.
Description	Displays the description of the split horizon group.
Instance Id	Displays the Instance identifier of the split horizon group.
Last Change	Displays the date and time of most recent change to the split horizon group.
Split Horizon Group	Displays the name of the split horizon group the SAP or Spoke SDP is associated.

Sample output (split horizon group)

```
*A:SASM>show>service# id 10 all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 10                  Vpn Id            : 0
Service Type      : VPLS
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 07/22/2011 11:06:02
Last Mgmt Change  : 07/22/2011 11:04:51
Admin State       : Up                  Oper State        : Up
MTU               : 1450
MTU Check         : Enabled
SAP Count         : 2                   SDP Bind Count    : 2
Snd Flush on Fail : Disabled
Uplink Type:      : MPLS

-------------------------------------------------------------------------------
Split Horizon Group specifics
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Split Horizon Group : test
-------------------------------------------------------------------------------
Description       : test
Instance Id       : 1                   Last Change       : 07/23/2011 11:40:50

-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
```

```
 Sdp Id 2:10  -(10.20.1.6)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 2:10                      Type              : Spoke
Split Horiz Grp   : (Not Specified)
VC Type           : VLAN                      VC Tag            : 10
Admin Path MTU    : 0                         Oper Path MTU     : 9186
Far End           : 10.20.1.6                 Delivery          : MPLS

Admin State       : Up                        Oper State        : Up
Acct. Pol         : None                      Collect Stats     : Disabled
Ingress Label     : 131063                    Egress Label      : 131067
Admin ControlWord : Preferred                 Oper ControlWord  : True
Last Status Change: 07/22/2011 11:07:26       Signaling         : TLDP
Last Mgmt Change  : 07/22/2011 11:04:51       Force Vlan-Vc     : Disabled
Flags             : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr  : 0                         Static MAC Addr   : 0

MAC Learning      : Enabled                   Discard Unkwn Srce: Disabled
BPDU Translation  : Disabled
L2PT Termination  : Disabled
MAC Pinning       : Disabled
MAC Pinning       : Disabled                  Block On Mesh Fail: False

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

Statistics          :
I. Fwd. Pkts.     : 0                         I. Fwd. Octs.     : 0
E. Fwd. Pkts.     : 1                         E. Fwd. Octets    : 98
Extra-Tag-Drop-Pkts: n/a                      Extra-Tag-Drop-Oc*: n/a

Associated LSP LIST :
Lsp Name          : toF
Admin State       : Up                        Oper State        : Up
-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Stp Admin State   : Up                        Stp Oper State    : Up
Core Connectivity : Down
Port Role         : Designated                Port State        : Forwarding
Port Number       : 2049                      Port Priority     : 128
Port Path Cost    : 10                        Auto Edge         : Enabled
Admin Edge        : Disabled                  Oper Edge         : True
Link Type         : Pt-pt                     BPDU Encap        : Dot1d
Root Guard        : Disabled                  Active Protocol   : Rstp
Last BPDU from    : N/A
Designated Bridge : This Bridge               Designated Port Id: 34817

Fwd Transitions   : 1                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd    : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd    : 0                         TCN BPDUs tx      : 0
TC bit BPDUs rcvd : 0                         TC bit BPDUs tx   : 0
RST BPDUs rcvd    : 0                         RST BPDUs tx      : 44265
-------------------------------------------------------------------------------
 Sdp Id 4:10  -(10.20.1.3)
```

```
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 4:10                      Type              : Spoke
Split Horiz Grp   : (Not Specified)
VC Type           : VLAN                      VC Tag            : 10
Admin Path MTU    : 0                         Oper Path MTU     : 9182
Far End           : 10.20.1.3                 Delivery          : MPLS

Admin State       : Up                        Oper State        : Up
Acct. Pol         : None                      Collect Stats     : Disabled
Ingress Label     : 131059                    Egress Label      : 131065
Admin ControlWord : Preferred                 Oper ControlWord  : True
Last Status Change: 07/22/2011 11:07:26       Signaling         : TLDP
Last Mgmt Change  : 07/22/2011 11:04:51       Force Vlan-Vc     : Disabled
Flags             : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr  : 0                         Static MAC Addr   : 0

MAC Learning      : Enabled                   Discard Unkwn Srce: Disabled
BPDU Translation  : Disabled
L2PT Termination  : Disabled
MAC Pinning       : Disabled
MAC Pinning       : Disabled                  Block On Mesh Fail: False

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

Statistics          :
I. Fwd. Pkts.     : 44285                     I. Fwd. Octs.     : 3852802
E. Fwd. Pkts.     : 0                         E. Fwd. Octets    : 0
Extra-Tag-Drop-Pkts: n/a                      Extra-Tag-Drop-Oc*: n/a

Associated LSP LIST :
Lsp Name          : toh2_facility
Admin State       : Up                        Oper State        : Up
Time Since Last Tr*: 01d00h37m


-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Stp Admin State   : Up                        Stp Oper State    : Up
Core Connectivity : Down
Port Role         : Root                      Port State        : Forwarding
Port Number       : 2050                      Port Priority     : 128
Port Path Cost    : 10                        Auto Edge         : Enabled
Admin Edge        : Disabled                  Oper Edge         : False
Link Type         : Pt-pt                     BPDU Encap        : Dot1d
Root Guard        : Disabled                  Active Protocol   : Rstp
Last BPDU from    : 80:01.00:25:ba:02:de:90
Designated Bridge : 80:01.00:25:ba:02:de:90   Designated Port Id: 34817

Fwd Transitions   : 1                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd    : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd    : 0                         TCN BPDUs tx      : 0
TC bit BPDUs rcvd : 2                         TC bit BPDUs tx   : 2
RST BPDUs rcvd    : 44284                     RST BPDUs tx      : 3
```

```
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/1/2
-------------------------------------------------------------------------------
Service Id         : 10
SAP                : 1/1/2                    Encap             : null
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Down
Flags              : PortOperDown
Last Status Change : 07/22/2011 11:04:50
Last Mgmt Change   : 07/23/2011 11:42:22
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
Admin MTU          : 1514                     Oper MTU          : 1514
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : None
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled

Acct. Pol          : None                     Collect Stats     : Disabled

-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                       Stp Oper State    : Up
Core Connectivity  : Down
Port Role          : Disabled                 Port State        : Discarding
Port Number        : 2051                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Rstp
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                      Designated Port   : 0

Forward transitions: 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
TC bit BPDUs rcvd  : 0                        TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                        Min Auth Interval : 15 minutes

-------------------------------------------------------------------------------
```

```
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1
-------------------------------------------------------------------------------
Aggregate Policer
-------------------------------------------------------------------------------
rate               : n/a                      burst             : n/a
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
Classifiers Allocated: 4                    Meters Allocated   : 2
Classifiers Used     : 2                    Meters Used        : 2
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                      Packets                Octets
Ingress Stats:        0                      0
Egress Stats:         0                      0
Ingress Drop Stats:   0                      0

Extra-Tag Drop Stats:   n/a                  n/a
-------------------------------------------------------------------------------
Sap per Meter stats
-------------------------------------------------------------------------------
                      Packets                Octets

Ingress Meter 1 (Unicast)
For. InProf           : 0                    0
For. OutProf          : 0                    0

Ingress Meter 11 (Multipoint)
For. InProf           : 0                    0
For. OutProf          : 0                    0

-------------------------------------------------------------------------------
SAP 1/1/7:10
-------------------------------------------------------------------------------
Service Id         : 10
SAP                : 1/1/7:10               Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                     Oper State        : Up
Flags              : None
Last Status Change : 07/22/2011 11:06:02
Last Mgmt Change   : 07/22/2011 11:04:51
Dot1Q Ethertype    : 0x8100                 QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Max Nbr of MAC Addr: No Limit               Total MAC Addr    : 2
Learned MAC Addr   : 0                      Static MAC Addr   : 2
Admin MTU          : 1518                   Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                    Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                    Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                    Egr IPv6 Fltr-Id  : n/a
tod-suite          : None
Mac Learning       : Enabled                Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled

Acct. Pol          : None                   Collect Stats     : Disabled
```

```
-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                       Stp Oper State    : Up
Core Connectivity  : Down
Port Role          : Designated               Port State        : Forwarding
Port Number        : 2048                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : True
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Rstp
Last BPDU from     : N/A
CIST Desig Bridge  : 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
TC bit BPDUs rcvd  : 0                        TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 44379
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                        Min Auth Interval : 15 minutes

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1
-------------------------------------------------------------------------------
Aggregate Policer
-------------------------------------------------------------------------------
rate               : n/a                      burst             : n/a
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
Classifiers Allocated: 4                     Meters Allocated   : 2
Classifiers Used     : 2                     Meters Used        : 2
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress Stats:          0                       0
Egress Stats:           1                       68
Ingress Drop Stats:     0                       0

Extra-Tag Drop Stats:   n/a                     n/a
-------------------------------------------------------------------------------
Sap per Meter stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Meter 1 (Unicast)
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Meter 11 (Multipoint)
For. InProf           : 0                       0
For. OutProf          : 0                       0
```

```
-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up                       Core Connectivity : Down
Stp Admin State    : Up                       Stp Oper State    : Up
Mode               : Rstp                     Vcp Active Prot.  : N/A

Bridge Id          : 80:02.00:25:ba:04:37:10  Bridge Instance Id: 2
Bridge Priority    : 32768                    Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 1d 00:38:51              Bridge Max Age    : 20
Top. Change Count  : 1                        Bridge Fwd Delay  : 15

Root Bridge        : 80:01.00:25:ba:02:de:90
Primary Bridge     : N/A

Root Path Cost     : 10                       Root Forward Delay: 15
Rcvd Hello Time    : 2                        Root Max Age      : 20
Root Priority      : 32769                    Root Port         : 2050

-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id        : 10                 Mac Move          : Disabled
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                Total Count       : 2
Learned Count     : 0                  Static Count      : 2
Remote Age        : 900                Local Age         : 300
High Watermark    : 95%                Low Watermark     : 90%
Mac Learning      : Enabled            Discard Unknown   : Disabled
Mac Aging         : Enabled            Relearn Only      : False

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
Endpoint name               : e1
Description                 : (Not Specified)
Revert time                 : 0
Act Hold Delay              : 0
Ignore Standby Signaling    : false
Suppress Standby Signaling  : false
Block On Mesh Fail          : false
Tx Active                   : none
Tx Active Up Time           : 0d 00:00:00
Revert Time Count Down      : N/A
Tx Active Change Count      : 0
Last Tx Active Change       : 07/22/2011 11:04:50
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
No members found.
===============================================================================
Endpoint name               : e2
Description                 : (Not Specified)
Revert time                 : 0
Act Hold Delay              : 0
Ignore Standby Signaling    : false
Suppress Standby Signaling  : false
Block On Mesh Fail          : false
Tx Active                   : none
```

```
Tx Active Up Time            : 0d 00:00:00
Revert Time Count Down       : N/A
Tx Active Change Count       : 0
Last Tx Active Change        : 07/22/2011 11:04:50
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
No members found.
===============================================================================
===============================================================================
```

arp

Syntax **arp** [*ip-address*] | [**mac** *ieee-address*] | [**sap** *sap-id*] | [**interface** *ip-int-name*]

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 interfaces ARP entry for easy lookup.

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

Default All IP addresses.

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

Default All MAC addresses.

sap *sap-id* **—** Displays SAP information for the specified SAP ID. See Common CLI Command Descriptions on page 987 for command syntax.

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

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

Values 1.0.0.0 — 223.255.255.255

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

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. Other — Local entries for the IP interfaces created.
Expiry	The age of the ARP entry.
Interface	The interface applied to the service.
SAP	The SAP ID.

Sample Output

```
*A:DUT-B# show service id 100 arp
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address         Type       Expiry    Interface         SAP
```

```
-------------------------------------------------------------------------------
192.168.1.2     00:00:01:00:00:01  Other     00h00m00s   HW                1/1/1:10*
195.168.1.1     32:67:01:01:00:03  Other     00h00m00s   to7x              1/1/3:10*
195.168.1.2     32:68:01:01:00:02  Dynamic   03h59m58s   to7x              1/1/3:10*
===============================================================================
*A:DUT-B#
```

base

Syntax **base**

Context show>service>id

Description This command displays basic information about this IES service.

Sample Output

```
*A:ALA-A# show service id 100 base
-------------------------------------------------------------
Service Basic Information
-------------------------------------------------------------
Service Id        : 100                Vpn Id            : 100
Service Type      : IES
Description       : Default Ies description for service id 100
Customer Id       : 1
Last Status Change: 08/29/2006 17:44:28
Last Mgmt Change  : 08/29/2006 17:44:28
Admin State       : Up                 Oper State        : Up
SAP Count         : 2
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                      Type         AdmMTU  OprMTU  Adm     Opr
-------------------------------------------------------------------------------
sap:1/1/3                       null         1514    1514    Up      Up
sap:1/1/4                       null         1514    1514    Up      Up
===============================================================================
*A:ALA-A#
```

interface

Syntax **interface** [*ip-address* | *ip-int-name*] [**detail**]

Context show>service>id

Description This command displays information for the IP interfaces associated with the IES service. If no optional parameters are specified, a summary of all IP interfaces associated to the service are displayed.

Parameters *ip-address* — The IP address of the interface for which to display information.

Values ipv4-address: a.b.c.d (host bits must be 0)

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

Values 32 characters maximum

detail — Displays detailed IP interface information.

Default IP interface summary output.

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

Label	Description
If Name	The name used to refer to the IES interface.
Type	Specifies the interface type.
IP-Address	Specifies the IP address/IP subnet/broadcast address of the interface.
Adm	The administrative state of the interface.
Opr	The operational state of the interface.
Admin State	The administrative state of the interface.
Oper State	The operational state of the interface.
IP Addr/mask	Specifies the IP address/IP subnet/broadcast address of the interface.
If Index	The index corresponding to this IES interface. The primary index is 1; all IES interfaces are defined in the base virtual router context.
If Type	Specifies the interface type.
SAP Id	Specifies the SAP's port ID.
SNTP B.Cast	Specifies whether SNTP broadcast client mode is enabled or disabled.
Arp Timeout	Specifies the timeout for an ARP entry learned on the interface.
MAC Address	Specifies the 48-bit IEEE 802.3 MAC address.
ICMP Mask Reply	Specifies whether ICMP mask reply is enabled or disabled.
Cflowd	Specifies whether Cflowd collection and analysis on the interface is enabled or disabled.
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         Type    Port/SapId
```

```
   IP-Address
-------------------------------------------------------------------------------
Sector A                          Up          Down/Down   IES     1/1/1.2.2
   -                                                              -
test                              Up          Down/Down   IES     1/1/2:0
   1.1.1.1/31                                                     n/a
   1.1.1.1/31                                                     n/a
   1.1.2.1/31                                                     n/a
test27                            Up          Up/--       IES Sub subscriber
   192.168.10.21/24                                               n/a
grp-if                            Up          Down/--     IES Grp 1/2/2
Interfaces : 4
===============================================================================
A:ALA-49#
A:ALA-49# show service id 88 interface
===============================================================================
Interface Table
===============================================================================
Interface-Name Adm Opr(v4/v6) Type Port/SapId
IP-Address PfxState
-------------------------------------------------------------------------------
Sector A Up Down/Down IES 1/1/1.2.2
- -
test Up Down/Down IES 1/1/2:0
1.1.1.1/31 n/a
1.1.1.1/31 n/a
1.1.2.1/31 n/a
test27 Up Up/-- IES Sub subscriber
192.168.10.21/24 n/a
grp-if Up Down/-- IES Grp 1/2/2
Interfaces : 4
===============================================================================
A:ALA-49#
```

VPRN Show Commands

egress-label

Syntax **egress-label** *start-label* **[***end-label***]**

Context show>service

Description Display services using the range of egress labels.

If only the mandatory *start-label* parameter is specified, only services using the specified label are displayed.

If both *start-label* and *end-label* parameters are specified, the services using the range of labels X where *start-label* <= X <= *end-label* are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *start-label* — The starting egress label value for which to display services using the label range. If only *egress-label1* is specified, services only using *egress-label1* are displayed.

Values 0 | 2048 — 131071

end-label — The ending egress label value for which to display services using the label range.

Default The *egress-label1* value.

Values 2049 — 131071

Output **Show Service Egress Command Output —** The following table describes show service egress label output fields.

Label	Description
Svc Id	The ID that identifies a service.
Sdp Id	The ID that identifies an SDP.
Type	Indicates whether the SDP binding is a spoke or a mesh.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
E. Lbl	The VC label used by this device to send packets to the far-end device in this service by the SDP.
Number of bindings found	The total number of SDP bindings that exist within the specified egress label range.

Sample Output

```
*A:ALA-12# show service egress-label 0 10000
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Id             Type I.Lbl        E.Lbl
-------------------------------------------------------------------------------
1          10:1               Mesh 0            0
1          20:1               Mesh 0            0
1          30:1               Mesh 0            0
1          100:1              Mesh 0            0
...
1          107:1              Mesh 0            0
1          108:1              Mesh 0            0
1          300:1              Mesh 0            0
1          301:1              Mesh 0            0
1          302:1              Mesh 0            0
1          400:1              Mesh 0            0
1          500:2              Spok 131070       2001
1          501:1              Mesh 131069       2000
100        300:100            Spok 0            0
200        301:200            Spok 0            0
300        302:300            Spok 0            0
400        400:400            Spok 0            0
-------------------------------------------------------------------------------
Number of Bindings Found : 23
===============================================================================
*A:ALA-12#
```

ingress-label

Syntax **ingress-label** *start-label* [*end-label*]

Context show>service

Description Display services using the range of ingress labels.

If only the mandatory *start-label* parameter is specified, only services using the specified label are displayed.

If both *start-label* and *end-label* parameters are specified, the services using the range of labels X where *start-label* <= X <= *end-label* are displayed.

Use the **show router** *vprn-service-id* **ldp bindings** command to display dynamic labels.

Parameters *start-label* — The starting ingress label value for which to display services using the label range. If only *start-label* is specified, services only using *start-label* are displayed.

Values 0, 2048 — 131071

end-label — The ending ingress label value for which to display services using the label range.

Default The *start-label* value.

Values 2048 — 131071

Output **Show Service Ingress-Label —** The following table describes show service ingress-label output fields:

Label	Description
Svc ID	The service identifier.
SDP Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
I.Lbl	The ingress label used by the far-end device to send packets to this device in this service by the SDP.
E.Lbl	The egress label used by this device to send packets to the far-end device in this service by the SDP.
Number of Bindings Found	The number of SDP bindings within the label range specified.

Sample Output

```
*A:ALA-12# show service ingress-label 0
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Id             Type I.Lbl        E.Lbl
-------------------------------------------------------------------------------
1          10:1               Mesh 0            0
1          20:1               Mesh 0            0
1          30:1               Mesh 0            0
1          50:1               Mesh 0            0
1          100:1              Mesh 0            0
1          101:1              Mesh 0            0
1          102:1              Mesh 0            0
1          103:1              Mesh 0            0
1          104:1              Mesh 0            0
1          105:1              Mesh 0            0
1          106:1              Mesh 0            0
1          107:1              Mesh 0            0
1          108:1              Mesh 0            0
1          300:1              Mesh 0            0
1          301:1              Mesh 0            0
1          302:1              Mesh 0            0
1          400:1              Mesh 0            0
100        300:100            Spok 0            0
200        301:200            Spok 0            0
300        302:300            Spok 0            0
400        400:400            Spok 0            0
-------------------------------------------------------------------------------
Number of Bindings Found : 21
-------------------------------------------------------------------------------
*A:ALA-12#
```

sap-using

Syntax **sap-using** [**sap** *sap-id*]
sap-using interface [*ip-address* | *ip-int-name*]
sap-using [**ingress** | **egress**] **filter** *filter-id*
sap-using [**ingress** | **egress**] **qos-policy** *qos-policy-id*

Context show>service

Description This command displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 987 for command syntax.

interface — Specifies matching SAPs with the specified IP interface.

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

Values 1.0.0.0 — 223.255.255.255

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

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

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

Values 1 — 65535

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.

Label	Description (Continued)
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#
```

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 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)
Opr MTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
IP address	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Adm Admin State	Specifies the state of the SDP.
Opr Oper State	Specifies the operating state of the SDP.
Flags	Specifies all the conditions that affect the operating status of this SDP.
Signal Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
Last Status Change	Specifies the time of the most recent operating status change to this SDP.
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SDP.
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Deliver Delivered	Specifies the type of delivery used by the SDP: MPLS.
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Hello Timeout	Specifies the number of seconds to wait for an SDP echo response message before declaring a timeout.
Unmatched Replies	Specifies the number of SDP unmatched message replies.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.

Label	Description (Continued)
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.

Sample Output

```
*A:ALA-12# show service sdp
===============================================================================
Services: Service Destination Points
===============================================================================
SdpId    Adm MTU   Opr MTU   IP address       Adm  Opr         Deliver Signal
-------------------------------------------------------------------------------
10       4462      4462      10.20.1.3        Up   Dn NotReady MPLS    TLDP
40       4462      1534      10.20.1.20       Up   Up          MPLS    TLDP
-------------------------------------------------------------------------------
Number of SDPs : 5
===============================================================================
*A:ALA-12#

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

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

Associated LSP LIST :
Lsp Name             : to-104
```

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

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**][**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.

sdp *sdp-id* **—** Displays only services bound to the specified SDP ID.

Default Services bound to any SDP ID.

Values 1 — 17407

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

Default Services associated with an customer.

Values 1 — 2147483647

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

Label	Description
Service Id	The service identifier.
Type	Specifies the service type configured for the service ID.
Adm	The desired state of the service.
Opr	The operating state of the service.
CustomerID	The ID of the customer who owns this service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this service.

Sample Output

```
*A:ALA-12# show service service-using customer 10
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
1            VPLS      Up     Up         10                09/05/2006 13:24:15
100          IES       Up     Up         10                09/05/2006 13:24:15
300          Epipe     Up     Up         10                09/05/2006 13:24:15
900          VPRN      Up     Up         2                 11/04/2006 04:55:12
-------------------------------------------------------------------------------
Matching Services : 4
===============================================================================
*A:ALA-12#

*A:ALA-12# show service service-using epipe
===============================================================================
Services [epipe]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
6            Epipe     Up     Up         6                 06/22/2006 23:05:58
7            Epipe     Up     Up         6                 06/22/2006 23:05:58
8            Epipe     Up     Up         3                 06/22/2006 23:05:58
103          Epipe     Up     Up         6                 06/22/2006 23:05:58
-------------------------------------------------------------------------------
Matching Services : 4
===============================================================================
*A:ALA-12#


A:de14# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
1            uVPLS     Up     Up         1                 10/26/2006 15:44:57
2            Epipe     Up     Down       1                 10/26/2006 15:44:57
10           mVPLS     Down   Down       1                 10/26/2006 15:44:57
```

```
11           mVPLS     Down   Down       1                 10/26/2006 15:44:57
100          mVPLS     Up     Up         1                 10/26/2006 15:44:57
101          mVPLS     Up     Up         1                 10/26/2006 15:44:57
102          mVPLS     Up     Up         1                 10/26/2006 15:44:57
999          uVPLS     Down   Down       1                 10/26/2006 16:14:33
-------------------------------------------------------------------------------
Matching Services : 8
-------------------------------------------------------------------------------
A:de14#
```

id

Syntax **id** *service-id* {**all** | **arp** | **base** | **fdb** | **labels** | **mfib** | **sap** | **sdp** | **split-horizon-group** | **stp**}

Context show>service

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

Parameters *service-id* — The unique service identification number that identifies the service in the service domain.

all — Display detailed information about the service.

arp — Display ARP entries for the service.

base — Display basic service information.

fdb — Display FDB entries.

interface — Display service interfaces.

labels — Display labels being used by this service.

sap — Display SAPs associated to the service.

sdp — Display SDPs associated with the service.

split-horizon-group — Display split horizon group information.

stp — Display STP information.

all

Syntax **all**

Context show>service>id

Description Displays detailed information for all aspects of the service.

Output **Sample Output**

```
*A:7210SAS>show>service>id# all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : Epipe
Description       : (Not Specified)
```

```
Customer Id       : 1
Last Status Change: 02/12/2002 23:51:07
Last Mgmt Change  : 02/12/2002 23:50:18
Admin State       : Up                  Oper State        : Up
SAP Count         : 2
Uplink Type:      : L2
SAP Type:         : Any                 Customer vlan:    : n/a
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/1/9:600.*
-------------------------------------------------------------------------------
Service Id         : 1
SAP                : 1/1/9:600.*             Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Up
Flags              : None
Last Status Change : 02/12/2002 23:51:06
Last Mgmt Change   : 02/12/2002 23:50:18
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100

Admin MTU          : 9212                    Oper MTU          : 9212
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
tod-suite          : None
Endpoint           : N/A

Acct. Pol          : None                    Collect Stats     : Disabled


-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : n/a
-------------------------------------------------------------------------------
Aggregate Policer
-------------------------------------------------------------------------------
rate               : n/a                     burst             : n/a
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
Classifiers Allocated: 2                   Meters Allocated   : 1
Classifiers Used     : 1                   Meters Used        : 1
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                       Packets                 Octets
Ingress Stats:         0                       0
Egress Stats:          26941105                18014193523

Extra-Tag Drop Stats:  n/a                     n/a

-------------------------------------------------------------------------------
SAP 1/1/12:90
-------------------------------------------------------------------------------
Service Id         : 1
SAP                : 1/1/12:90               Encap             : q-tag
```

```
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Up
Flags              : None
Last Status Change : 02/12/2002 23:51:07
Last Mgmt Change   : 02/13/2002 00:05:46
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100
Loopback Mode      : Internal                 No-svc-port used  : 1/1/25
Loopback Src Addr  : 00:00:01:00:02:00
Loopback Dst Addr  : 00:00:01:00:03:00

Admin MTU          : 1518                     Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : None
Endpoint           : N/A

Acct. Pol          : None                     Collect Stats     : Disabled


-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1
-------------------------------------------------------------------------------
Aggregate Policer
-------------------------------------------------------------------------------
rate               : n/a                      burst             : n/a
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
Classifiers Allocated: 2                    Meters Allocated   : 1
Classifiers Used     : 1                    Meters Used        : 1
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                      Packets                 Octets
Ingress Stats:        26940595                18013850572
Egress Stats:         0                       0
Ingress Drop Stats:   0                       0

Extra-Tag Drop Stats:   n/a                   n/a
-------------------------------------------------------------------------------
Sap per Meter stats (in/out counter mode)
-------------------------------------------------------------------------------
                      Packets                 Octets

Ingress Meter 1
For. InProf          : 8                      4265
For. OutProf         : 26941156               18014224039

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:7210SAS>show>service>id#
```

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

Label	Description
Service Detailed Information	
Service Id	The service identifier.
VPN Id	The number which identifies the VPN.
Customer Id	The customer identifier.
Last Status Change	The date and time of the most recent change in the administrative or operating status of the service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Admin State	The current administrative state.
Oper State	The current operational state.
Route Dist.	Displays the route distribution number.
AS Number	Displays the autonomous system number.
Router Id	Displays the router ID for this service.
Auto Bind	Specifies the automatic binding type for the SDP assigned to this service.
Vrf Target	Specifies the VRF target applied to this service.
Vrf Import	Specifies the VRF import policy applied to this service.
Vrf Export	Specifies the VRF export policy applied to this service.
Description	Generic information about the service.
SAP Count	The number of SAPs specified for this service.
SDP Bind Count	The number of SDPs bound to this service.
Split Horizon Group	Name of the split horizon group for this service.
Description	Description of the split horizon group.
Last Changed	The date and time of the most recent management-initiated change to this split horizon group.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.
Type	Indicates whether this Service SDP binding is a spoke or a mesh.
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.

Label	Description
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the keepalive protocol.
Oper State	The current status of the keepalive protocol.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Number of SDPs	The total number SDPs applied to this service ID.
Service Access Points	

Label	Description
Service Id	The service identifier.
Port Id	The ID of the access port where this SAP is defined.
Description	Generic information about the SAP.
Admin State	The desired state of the SAP.
Oper State	The operating state of the SAP.
Last Changed	The date and time of the last change.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The SAP ingress QoS policy ID.
Egress qos-policy	The SAP egress QoS policy ID. This command is applicable only to 7210 SAS X.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
Spoke SDPs	
Managed by Service	Specifies the service-id of the management VPLS managing this spoke SDP.
Managed by Spoke	Specifies the sap-id inside the management VPLS managing this spoke SDP.
Prune state	Specifies the STP state inherited from the management VPLS.
Peer Pw Bits	Indicates the bits set by the LDP peer when there is a fault on its side of the pseudowire. LAC failures occur on the SAP that has been configured on the pipe service, PSN bits are set by SDP-binding failures on the pipe service. The pwNotForwarding bit is set when none of the above failures apply, such as an MTU mismatch failure. This value is only applicable if the peer is using the pseudowire status signalling method to indicate faults. pwNotForwarding — Pseudowire not forwarding lacIngressFault Local — Attachment circuit RX fault lacEgresssFault Local — Attachment circuit TX fault psnIngressFault Local — PSN-facing PW RX fault psnEgressFault Local — PSN-facing PW TX fault pwFwdingStandby — Pseudowire in standby mode

Label	Description
Max IPv4 Routes	Maximum IPv4 routes configured for use with the service.
Last Changed	The date and time of the most recent management-initiated change.
Dot1Q Ethertype	The Dot1q ethertype in use by the SAP.
Ingr IP Fltr-Id	The policy ID of the IP filter applied at ingress.
Ingr Mac Fltr-Id	The policy ID of the MAC filter applied at ingress.
Egr IP Fltr-Id	The policy ID of the IP filter applied at egress.
Egr Mac Fltr-Id	The policy ID of the MAC filter applied at egress.
tod-suite	The TOD suite applied for use by this SAP.
rate	Specifies the SAP aggregate rate configured for the aggregate policer/ meter used by this SAP.
burst	Specifies the burst to be used with SAP aggregate policer/meter used by this SAP.
Classifiers Allo-cated	Number of SAP ingress QoS resources allocated for use by this SAP.
Classifiers Used	Number of SAP ingress QoS resources in use by this SAP.
Meters Allocated	Number of SAP ingress meter resources allocated for use by this SAP. This is set to half the number of classifiers allocated to this SAP.
Meters Used	Number of SAP ingress meters in use.
Ingress Stats	The number of received packets/octets for this SAP.
Egress Stats	The number of packets/octets forwarded out of this SAP.
Ingress Drop Stats	Number of packets/octets dropped by the system.
Extra-Tag Drop Stats	Number of packets received with the count of VLAN tags exceeding the count of VLAN tags implied by the SAP encapsulation.
Ingress Meter 1	The index of the ingress QoS meter of this SAP.
For. InProf	Number of in-profile packets/octets received on this SAP.
For. OutProf	Number of out-of-profile packets/octets received on this SAP.
If Name	IP interface name assigned by user.
Protocols	Protocols enabled for use on this interface.
Oper (v4/v6)	Operational status of this interface for IPv4 and IPv6.
IP Addr/mask	IPv4 address and Mask assigned to this interface.
Address Type	Whether the address is a primary or secondary address.
Broadcast Address	Type of broadcast address used. It can be host-ones or all-ones.

Label	Description
If Index	The interface Index assigned by the system. It is used with SNMP IfTable.
Virt. If Index	The interface index assigned by the system. It is used with SNMP.
Last Oper Chg	Timestamp associated with the last operational change.
Global If Index	This is the system wide Interface index allotted by the system.
If Type	Network — The IP interface is a network/core IP interface. Service — The IP interface is a service IP interface.
SNTP B.Cast	Specifies whether SNTP broadcast client mode is enabled or disabled.
Arp Timeout	Specifies the timeout for an ARP entry learned on the interface.
IP Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the port to the far-end router, without requiring the packet to be fragmented.
LdpSyncTimer	Specifies the value used for IGP-LDP synchronization.
Redirects	Specifies the rate for ICMP redirect messages.
Unreachables	Specifies the rate for ICMP unreachable messages.
TTL Expired	Specifies the rate for ICMP TTL messages.
MAC Address	Specifies the 48-bit IEEE 802.3 MAC address.

authentication

Syntax **authentication**

Context show>service>id

Description This command enables the context to display subscriber authentication information.

statistics

Syntax **statistics** [**policy** *name*] [**sap** *sap-id*]

Context show>service>id>authentication

Description This command displays session authentication statistics for this service.

Parameters **policy** *name* — Specifies the subscriber authentication policy statistics to display.

sap *sap-id* — Specifies the SAP ID statistics to display. See Common CLI Command Descriptions on page 987 for command syntax.

Sample Output

```
*A:ALA-1# show service id 11 authentication statistics
===============================================================
Authentication statistics
===============================================================
Interface / SAP                Authentication  Authentication
                               Successful      Failed
---------------------------------------------------------------
abc-11-90.1.0.254              1582            3
---------------------------------------------------------------
Number of entries: 1
===============================================================
*A:ALA-1#
```

arp

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

Context show>service>id

Description Displays the ARP table for the IES instance.

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

Default All IP addresses.

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

Default All MAC addresses.

sap *sap-id* **—** Displays SAP information for the specified SAP ID. See Common CLI Command Descriptions on page 987 for command syntax.

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

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

Values 1.0.0.0 — 223.255.255.255

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

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

Label	Description
Service ID	The service ID number.
MAC	The specified MAC address
Source-Identifier	The location the MAC is defined.

Type	Static – FDB entries created by management. Learned – Dynamic entries created by the learning process. OAM – Entries created by the OAM process.
Age	The time elapsed since the service was enabled.
Interface	The interface applied to the service.
Port	The port where the SAP is applied.

Sample Output

```
*A:ALA-12# show service id 2 arp
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address       Type    Age       Interface         Port
-------------------------------------------------------------------------------
190.11.1.1      00:03:fa:00:08:22 Other   00:00:00  ies-100-190.11.1  1/1/11:0
===============================================================================
*A:ALA-12#
```

base

Syntax **base**

Context show>service>id

Description Displays basic information about the service ID including service type, description, SAPs and SDPs.

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

Label	Description
Service Id	The service identifier.
Vpn Id	Specifies the VPN ID assigned to the service.
Service Type	Specifies the type of service.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Adm	The desired state of the service.
Oper	The operating state of the service.
Mtu	The largest frame size (in octets) that the service can handle.

Label	Description
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#
```

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. See Common CLI Command Descriptions on page 973 for command syntax.

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.

Label	Description
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#
```

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.
ICMP Mask Reply	Specifies whether ICMP mask reply is enabled or disabled.
ICMP Details	
Redirects	Specifies the rate for ICMP redirect messages.
Unreachables	Specifies the rate for ICMP unreachable messages.
TTL Expired	Specifies the rate for ICMP TTL messages.

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       : 71/1/1.2.2
TOS Marking  : Untrusted                         If Type          : IES
SNTP B.Cast  : False                             IES ID           : 88
MAC Address  : Not configured.                   Arp Timeout      : 14400
IP MTU       : 1500                              ICMP Mask Reply  : True
Arp Populate : Disabled
Cflowd       : None

Proxy ARP Details
Proxy ARP    : Enabled                           Local Proxy ARP  : Disabled
Policies     : ProxyARP

DHCP Details
Admin State  : Up                                Lease Populate   : 0
Action       : Keep                              Trusted          : Disabled
ICMP Details
Redirects    : Number - 100                      Time (seconds)   - 10
Unreachables : Number - 100                      Time (seconds)   - 10
TTL Expired  : Number - 100                      Time (seconds)   - 10
-------------------------------------------------------------------------------
Interface
-------------------------------------------------------------------------------
If Name      : test
Admin State  : Up                                Oper State       : Down
Protocols    : None
IP Addr/mask : Not Assigned
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Description  :
If Index     : 27                                Virt. If Index   : 27
SAP Id       : 101/1/2:0
TOS Marking  : Untrusted                         If Type          : IES
SNTP B.Cast  : False                             IES ID           : 88
MAC Address  : Not configured.                   Arp Timeout      : 14400
Arp Populate : Disabled

Proxy ARP Details
Proxy ARP    : Disabled                          Local Proxy ARP  : Disabled

ICMP Details
Redirects    : Number - 100                      Time (seconds)   - 10
Unreachables : Number - 100                      Time (seconds)   - 10
```

```
TTL Expired  : Number - 100                    Time (seconds)   - 10
-------------------------------------------------------------------------------
Interfaces : 2
===============================================================================
A:ALA-49#
```

sap

Syntax **sap** *sap-id* [**detail**]]

Context show>service>id

Description Displays information for the SAPs associated with the service.

If no optional parameters are specified, a summary of all associated SAPs is displayed.

Parameters *sap-id* — The ID that displays SAPs for the service. See Common CLI Command Descriptions on page 987 for command syntax.

detail — Displays detailed information for the SAP.

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

Sample Output

Label	Description
Service Id	The service identifier.
SAP	The SAP and qtag.
Encap	The encapsulation type of the SAP.
Ethertype	Specifies an Ethernet type II Ethertype value.
Admin State	The administrative state of the SAP.
Oper State	The operating state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch,RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
Last Status Change	Specifies the time of the most recent operating status change to this SAP
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SAP.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.
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.

Label	Description (Continued)
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP.
Egress qos-policy	The egress QoS policy ID assigned to the SAP.
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP.
Egress Filter-Id	The egress filter policy ID assigned to the SAP.
Acct. Pol	The accounting policy ID assigned to the SAP.
Collect Stats	Specifies whether collect stats is enabled.
Dropped	The number of packets and octets dropped due to SAP state, ingress MAC or IP filter, same segment discard, bad checksum, etc.
Off. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Off. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Off. Uncolor	The number of uncolored packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Dro. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to: MBS exceeded, buffer pool limit exceeded, etc.
Dro. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to: MBS exceeded, buffer pool limit exceeded, etc.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded by the ingress Qchip.
For. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
Dro. InProf	The number of in-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
Dro. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded by the egress Qchip.
For. OutProf	The number of out-of-profile packets and octets (rate above CIR) forwarded by the egress Qchip.
Ingress TD Profile	The profile ID applied to the ingress SAP.
Egress TD Profile	The profile ID applied to the egress SAP.

Label	Description (Continued)
Alarm Cell Handling	The indication that OAM cells are being processed.
AAL-5 Encap	The AAL-5 encapsulation type.

```
*A:ALA-12# show service id 321 sap 1/1/4:0
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 321
SAP                : 1/1/4:0                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 1/1/4:0 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 321
SAP                : 1/1/4:0                Encap             : q-tag
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100

Admin State        : Up                      Oper State        : Down
Flags              : PortOperDown
                     SapIngressQoSMismatch
Last Status Change : 02/03/2007 12:58:37
Last Mgmt Change   : 02/03/2007 12:59:10
Admin MTU          : 1518                    Oper MTU          : 1518
Ingress qos-policy : 100                     Egress qos-policy : 1
Ingress Filter-Id  : n/a                     Egress Filter-Id  : n/a
Multi Svc Site     : None
Acct. Pol          : None                    Collect Stats     : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                       Packets                 Octets
Forwarding Engine Stats
Dropped              : 0                       0
Off. HiPrio          : 0                       0
Off. LowPrio         : 0                       0
Off. Uncolor         : 0                       0
Queueing Stats(Egress QoS Policy 1)
Dro. InProf          : 0                       0
Dro. OutProf         : 0                       0
For. InProf          : 0                       0
For. OutProf         : 0                       0
-------------------------------------------------------------------------------
===============================================================================
```

```
*A:ALA-12#
```

sdp

Syntax **sdp** [*sdp-id* | **far-end** *ip-addr*] [**detail**]

Context show>service>id

Description Displays information for the SDPs associated with the service. If no optional parameters are specified, a summary of all associated SDPs is displayed.

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

Default All SDPs.

Values 1 — 17407

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

Default SDPs with any far-end IP address.

detail — Displays detailed SDP information.

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

Label	Description
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
Split Horizon Group	Name of the split horizon group that the SDP belongs to.
VC Type	Displays the VC type: ether or vlan.
VC Tag	Displays the explicit dot1Q value used when encapsulating to the SDP far end.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case.)
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.

Label	Description (Continued)
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the keepalive process.
Oper State	he operational state of the keepalive process.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts.	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field.

Sample Output

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

Admin State         : Up                        Oper State        : Up
Acct. Pol           : None                      Collect Stats     : Disabled
Ingress Label       : 2048                      Egress Label      : 2048
Ing mac Fltr        : n/a                       Egr mac Fltr      : n/a
Ing ip Fltr         : n/a                       Egr ip Fltr       : n/a
Ing ipv6 Fltr       : n/a                       Egr ipv6 Fltr     : n/a
Admin ControlWord   : Not Preferred             Oper ControlWord  : False
```

```
Last Status Change : 05/31/2007 00:45:43      Signaling         : None
Last Mgmt Change   : 05/31/2007 00:45:43
Class Fwding State : Up
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0

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

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

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

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

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

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

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

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

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

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

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

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

Lsp Name           : A_B_10
Admin State        : Up                        Oper State        : Up
Time Since Last Tr*: 00h26m34s
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : enabled
Default LSP        : A_B_10                    Multicast LSP     : A_B_9
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
af                  A_B_3
be                  A_B_1
ef                  A_B_6
h1                  A_B_7
h2                  A_B_5
l1                  A_B_4
l2                  A_B_2
nc                  A_B_8
===============================================================================
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Up                        Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                       Port State        : Forwarding
Port Number        : 2049                      Port Priority     : 128
Port Path Cost     : 10                        Auto Edge         : Enabled
Admin Edge         : Disabled                  Oper Edge         : N/A
Link Type          : Pt-pt                     BPDU Encap        : Dot1d
Root Guard         : Disabled                  Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                       Designated Port Id: 0

Fwd Transitions    : 0                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                         TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                         RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
* indicates that the corresponding row element may have been truncated.
-------------------------------------------------------------------------------
A:Dut-A#
```

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.

Label	Description (Continued)
MAC Address	The MAC address of the ARP entry.
Expiry	The age of the ARP entry.
Type	Dyn – The ARP entry is a dynamic ARP entry.
	Inv – The ARP entry is an inactive static ARP entry (invalid).
	Oth – The ARP entry is a local or system ARP entry.
	Sta – The ARP entry is an active static ARP entry.
Interface	The IP interface name associated with the ARP entry.
No. of ARP Entries	The number of ARP entries displayed in the list.

Sample Output

```
*A:ALA-12# show router 3 arp
===============================================================================
ARP Table (Service: 3)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
10.10.10.103    04:67:ff:00:00:01 00h00m00s Oth    system
10.10.4.3       00:00:00:00:00:00 00h00m00s Oth    ALA-1-2
10.10.5.3       00:00:00:00:00:00 00h00m00s Oth    ALA-1-3
10.10.7.3       00:00:00:00:00:00 00h00m00s Oth    ALA-1-5
10.10.0.16      00:00:00:00:00:00 00h00m00s Oth    bozo
10.10.3.3       00:00:00:00:00:00 00h00m00s Oth    gizmo
10.10.2.3       00:00:00:00:00:00 00h00m00s Oth    hobo
10.10.1.17      00:00:00:00:00:00 00h00m00s Oth    int-cflowd
10.0.0.92       00:00:00:00:00:00 04h00m00s Dyn    to-104
10.0.0.103      04:67:01:01:00:01 00h00m00s Oth[I] to-104
10.0.0.104      04:68:01:01:00:01 03h59m49s Dyn[I] to-104
10.10.36.2      00:00:00:00:00:00 00h00m00s Oth    tuesday
192.168.2.98    00:03:47:c8:b4:86 00h14m37s Dyn[I] management
192.168.2.103   00:03:47:dc:98:1d 00h00m00s Oth[I] management
-------------------------------------------------------------------------------
No. of ARP Entries: 14
===============================================================================
*A:ALA-12#


*A:ALA-12# show router 3 arp 10.10.0.3
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
10.10.0.3       04:5d:ff:00:00:00 00:00:00  Oth    system
===============================================================================
*A:ALA-12#


*A:ALA-12# show router 3 arp to-ser1
===============================================================================
```

```
ARP Table
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
10.10.13.1      04:5b:01:01:00:02 03:53:09  Dyn    to-ser1
===============================================================================
*A:ALA-12#
```

damping

Syntax **damping** [*ip-prefix/mask* | *ip-address*] [**detail**]
damping [*damp-type*] [**detail**]

Context show>router>bgp

Description This command displays BGP routes with have been dampened due to route flapping. This command can be entered with or without a route parameter.

When the keyword **detail** is included, more detailed information displays.

When only the command is entered (without any parameters included except **detail**), then all dampened routes are listed.

When a parameter is specified, then the matching route or routes are listed.

When a **decayed**, **history,** or **suppressed** keyword is specified, only those types of dampened routes are listed.

Parameters *ip-prefix/mask* — Displays damping information for the specified IP prefix and mask length.

ip-address — Displays damping entry for the best match route for the specified IP address.

damp-type — Displays damping type for the specified IP address.

decayed — Displays damping entries that are decayed but are not suppressed.

history — Displays damping entries that are withdrawn but have history.

suppressed — Displays damping entries suppressed because of route damping.

detail — Displays detailed information.

Output **Show Damping Output Fields —** The following table describes BGP damping output fields:

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured or inherited local AS for the specified peer group. If not configured, then it is the same value as the AS.
Network	Route IP prefix and mask length for the route.
Flag(s)	Legend: Status codes: u- used, s-suppressed, h-history, d-decayed, *-valid. If a * is not present, then the status is invalid. Origin codes: i-IGP, e-EGP, ?-incomplete, >-best
Network	The IP prefix and mask length for the route.
From	The originator ID path attribute value.
Reuse time	The time when a suppressed route can be used again.
AS Path	The BGP AS path for the route.

Label	Description (Continued)
Peer	The router ID of the advertising router.
NextHop	BGP nexthop for the route.
Peer AS	The autonomous system number of the advertising router.
Peer Router-Id	The router ID of the advertising router.
Local Pref	BGP local preference path attribute for the route.
Age	The time elapsed since the service was enabled.
Last update	The time when BGP was updated last in second/minute/hour (SS:MM:HH) format.
FOM Present	The current Figure of Merit (FOM) value.
Number of Flaps	The number of flaps in the neighbor connection.
Reuse time	The time when the route can be reused.
Path	The BGP AS path for the route.
Applied Policy	The applied route policy name.

Sample Output

```
*A:ALA-12# show router 3 bgp damping
===============================================================================
  BGP Router ID : 10.0.0.14         AS : 65206   Local AS : 65206
===============================================================================
  Legend -
  Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
  Origin codes  : i - IGP, e - EGP, ? - incomplete,  - best
===============================================================================
BGP Damped Routes
===============================================================================
Flag  Network            From            Reuse       AS-Path
-------------------------------------------------------------------------------
ud*i  12.149.7.0/24      10.0.28.1       00h00m00s   60203 65001 19855 3356
                                                      1239  22406
si    24.155.6.0/23      10.0.28.1       00h43m41s   60203 65001 19855 3356
                                                      2914  7459
si    24.155.8.0/22      10.0.28.1       00h38m31s   60203 65001 19855 3356
                                                      2914  7459
si    24.155.12.0/22     10.0.28.1       00h35m41s   60203 65001 19855 3356
                                                      2914  7459
si    24.155.22.0/23     10.0.28.1       00h35m41s   60203 65001 19855 3356
                                                      2914  7459
si    24.155.24.0/22     10.0.28.1       00h35m41s   60203 65001 19855 3356
                                                      2914  7459
si    24.155.28.0/22     10.0.28.1       00h34m31s   60203 65001 19855 3356
                                                      2914  7459
si    24.155.40.0/21     10.0.28.1       00h28m24s   60203 65001 19855 3356
                                                      7911  7459
si    24.155.48.0/20     10.0.28.1       00h28m24s   60203 65001 19855 3356
                                                      7911  7459
ud*i  61.8.140.0/24      10.0.28.1       00h00m00s   60203 65001 19855 3356
```

```
                                                    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:

Sample Output

Label	Description
Group	BGP group name
Group Type	No Type — Peer type not configured.
	External — Peer type configured as external BGP peers.
	Internal — Peer type configured as internal BGP peers.
State	Disabled — The BGP peer group has been operationally disabled.
	Down — The BGP peer group is operationally inactive.
	Up — The BGP peer group is operationally active.
Peer AS	The configured or inherited peer AS for the specified peer group.
Local AS	The configured or inherited local AS for the specified peer group.
Local Address	The configured or inherited local address for originating peering for the specified peer group.
Loop Detect	The configured or inherited loop detect setting for the specified peer group.
Connect Retry	The configured or inherited connect retry timer value.

Label	Description (Continued)
	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.
Multipath	The configured or inherited multipath value, determining the maximum number of ECMP routes BGP can advertise to the RTM.
Prefix Limit	No Limit – No route limit assigned to the BGP peer group. 1 - 4294967295 – The maximum number of routes BGP can learn from a peer.
Passive	Disabled – BGP attempts to establish BGP connections with neighbors in the specified peer group. Enabled – BGP will not actively attempt to establish BGP connections with neighbors in the specified peer group.
Next Hop Self	Disabled – BGP is not configured to send only its own IP address as the BGP nexthop in route updates to neighbors in the peer group. Enabled – BGP sends only its own IP address as the BGP nexthop in route updates to neighbors in the specified peer group.
Aggregator ID 0	Disabled – BGP is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group. Enabled – BGP is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group.
Remove Private	Disabled – BGP will not remove all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group. Enabled – BGP removes all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group.
Damping	Disabled – The peer group is configured not to dampen route flaps.

Label	Description (Continued)
	Enabled — The peer group is configured to dampen route flaps.
Export Policy	The configured export policies for the peer group.
Import Policy	The configured import policies for the peer group.
Hold Time	The configured hold time setting.
Keep Alive	The configured keepalive setting.
Cluster Id	None — No cluster ID has been configured.
Client Reflect	Disabled — The BGP route reflector will not reflect routes to this neighbor.
	Enabled — The BGP route reflector is configured to reflect routes to this neighbor.
NLRI	The type of NLRI information that the specified peer group can accept.
	Unicast — IPv4 unicast routing information can be carried.
Preference	The configured route preference value for the peer group.
List of Peers	A list of BGP peers configured under the peer group.
Total Peers	The total number of peers configured under the peer group.
Established	The total number of peers that are in an established state.

```
*A:ALA-12# show router 3 bgp group
===============================================================================
BGP Groups
===============================================================================
Group            : To_AS_40000
-------------------------------------------------------------------------------
Description      : Not Available
Group Type       : No Type               State            : Up
Peer AS          : 40000                 Local AS         : 65206
Local Address    : n/a                   Loop Detect      : Ignore
Export Policy    : direct2bgp
Hold Time        : 90                    Keep Alive       : 30
Cluster Id       : None                  Client Reflect   : Enabled
NLRI             : Unicast               Preference       : 170

List of Peers
- 10.0.0.1       : To_Jukebox
- 10.0.0.12      : Not Available
- 10.0.0.13      : Not Available
- 10.0.0.14      : To_ALA-1
- 10.0.0.15      : To_H-215
Total Peers      : 5                     Established      : 2
===============================================================================
*A:ALA-12#
```

neighbor

Syntax **neighbor** [*ip-address* [[**family** *family*] *filter1*]]
neighbor [*as-number* [[**family** *family*] filter2]]

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.

Label	Description (Continued)
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.
	Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION.
Last Event	start — BGP has initialized the BGP neighbor.
	stop — BGP has disabled the BGP neighbor.
	open — BGP transport connection opened.
	close — BGP transport connection closed.
	openFail — BGP transport connection failed to open.
	error — BGP transport connection error.

Label	Description (Continued)
	`connectRetry` — Connect retry timer expired. `holdTime` — Hold time timer expired. `keepAlive` — Keepalive timer expired. `recvOpen` — Receive an OPEN message. `revKeepalive` — Receive an KEEPALIVE message. `recvUpdate` — Receive an UPDATE message. `recvNotify` — Receive an NOTIFICATION message. `None` — No events have occurred.
`Last Error`	Displays the last BGP error and sub-code to occur on the BGP neighbor.
`Connect Retry`	The configured or inherited connect retry timer value.
`Local Pref.`	The configured or inherited local preference value.
`Min Route Advt.`	The minimum amount of time that must pass between route updates for the same IP prefix.
`Min AS Originate`	The minimum amount of time that must pass between updates for a route originated by the local router.
`Multihop`	The maximum number of router hops a BGP connection can traverse.
`Multipath`	The configured or inherited multipath value, determining the maximum number of ECMP routes BGP can advertise to the RTM.
`Damping`	`Disabled` — BGP neighbor is configured not to dampen route flaps. `Enabled` — BGP neighbor is configured to dampen route flaps.
`Loop Detect`	`Ignore` — The BGP neighbor is configured to ignore routes with an AS loop. `Drop` — The BGP neighbor is configured to drop the BGP peering if an AS loop is detected. `Off` — AS loop detection is disabled for the neighbor.
`MED Out`	The configured or inherited MED value assigned to advertised routes without a MED attribute.
`Authentication`	`None` — No authentication is configured. `MD5` — MD5 authentication is configured.

Label	Description (Continued)
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	The configured route reflector cluster ID. None — No cluster ID has been configured
Client Reflect	Disabled — The BGP route reflector is configured not to reflect routes to this neighbor. Enabled — The BGP route reflector is configured to reflect routes to this neighbor.
Preference	The configured route preference value for the peer group.
Num of Flaps	The number of flaps in the neighbor connection.
Recd. Prefixes	The number of routes received from the BGP neighbor.
Active Prefixes	The number of routes received from the BGP neighbor and active in the forwarding table.

Label	Description (Continued)
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.
Flag	u – used
	s – suppressed
	h – history

Label	Description (Continued)
	d – decayed
	* – valid
	i – igp
	? – incomplete
	> – best
Network	Route IP prefix and mask length for the route.
Next Hop	BGP nexthop for the route.
LocalPref	BGP local preference path attribute for the route.
MED	BGP Multi-Exit Discriminator (MED) path attribute for the route.
AS Path	The BGP AS path for the route.

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

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	The aggregator AS value. none – No aggregator AS attributes are present.
Aggregator	The aggregator attribute value. none – no Aggregator attributes are present.
Atomic Aggr.	Atomic – The atomic aggregator flag is set. Not Atomic – The atomic aggregator flag is not set.
MED	The MED metric value. none – No MED metric is present.
Community	The BGP community attribute list.
Cluster	The route reflector cluster list.

Label	Description
Originator Id	The originator ID path attribute value. none — The originator ID attribute is not present.
Peer Router Id	The router ID of the advertising router.
AS-Path	The BGP AS path attribute.
VPRN Imported	Displays the VPRNs where a particular BGP-VPN received route has been imported and installed.

Sample Output

```
*A:ALA-12>config>router>bgp# show router 3 bgp routes family ipv4
===============================================================================
 BGP Router ID : 10.10.10.103      AS : 200     Local AS : 200
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP Routes
===============================================================================
Flag  Network                           Nexthop         LocalPref  MED
      VPN Label                         As-Path
-------------------------------------------------------------------------------
No Matching Entries Found
===============================================================================
*A:ALA-12>config>router>bgp#

A:SR-12# show router bgp routes 100.0.0.0/31 hunt
===============================================================================
 BGP Router ID : 10.20.1.1   AS : 100Local AS : 100
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP Routes
===============================================================================
RIB In Entries
-------------------------------------------------------------------------------
Network        : 100.0.0.0/31
Nexthop        : 10.20.1.2
Route Dist.    : 10.20.1.2:1         VPN Label     : 131070
From       : 10.20.1.2
Res. Nexthop   : 10.10.1.2
Local Pref.    : 100                 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 Confedera-tions	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.

Label	Description
Total BGP Active Routes	The total number of BGP routes used in the forwarding table.
Total BGP Routes	The total number of BGP routes learned from BGP peers.
Total BGP Paths	The total number of unique sets of BGP path attributes learned from BGP peers.
Total Path Memory	Total amount of memory used to store the path attributes.
Total Suppressed Routes	Total number of suppressed routes due to route damping.
Total History Routes	Total number of routes with history due to route damping.
Total Decayed Routes	Total number of decayed routes due to route damping.
Neighbor	BGP neighbor address.
AS (Neighbor)	BGP neighbor autonomous system number.
PktRcvd	Total number of packets received from the BGP neighbor.
PktSent	Total number of packets sent to the BGP neighbor.
InQ	The number of BGP messages to be processed.
OutQ	The number of BGP messages to be transmitted.
Up/Down	The amount of time that the BGP neighbor has either been established or not established depending on its current state.
State|Recv/Actv/Sent	The BGP neighbor's current state (if not established) or the number of received routes, active routes and sent routes (if established).

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#
```

interface

Syntax **interface** [{[<ip-address|ip-int-name>][**detail**]}|**summary**]

Context show>router

Description This command displays the router IP interface table sorted by interface index.

Parameters *ip-address* — Only displays the interface information associated with the specified IP address.

ip-int-name — Only displays the interface information associated with the specified IP interface name.

detail — Displays detailed IP interface information.

summary — Displays summary IP interface information for the router.

— 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:7210SAS>show>router interface i1 detail

```
===============================================================================
Interface Table (Router: Base)
===============================================================================

-------------------------------------------------------------------------------
```

```
Interface
-------------------------------------------------------------------------------
If Name      : i1
Admin State  : Up                                Oper (v4/v6)     : Down/--
Protocols    : None

IP Addr/mask : Not Assigned
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Description  : (Not Specified)
If Index     : 2                                 Virt. If Index   : 2
Last Oper Chg: 03/07/2001 01:47:29               Global If Index  : 127
Port Id      : 1/1/1
TOS Marking  : Trusted                           If Type          : Network
Egress Filter: none                              Ingress Filter   : none
Egr IPv6 Flt : none                              Ingr IPv6 Flt    : none
SNTP B.Cast  : False                             QoS Policy       : 2
Queue-group  : None
MAC Address  : 00:25:ba:0d:27:32                 Arp Timeout      : 14400
IP Oper MTU  : 9198
LdpSyncTimer : None                              Strip-Label      : Disabled
uRPF Chk     : disabled
uRPF Fail By*: 0                                 uRPF Chk Fail Pk*: 0

ICMP Details
Redirects    : Number - 100                      Time (seconds)   - 10
Unreachables : Number - 100                      Time (seconds)   - 10
TTL Expired  : Number - 100                      Time (seconds)   - 10

===============================================================================
Meter Statistics
===============================================================================

-------------------------------------------------------------------------------
                        Packets                  Octets
-------------------------------------------------------------------------------
Ingress Meter 1 (Unicast)
For. InProf           : 0                        0
For. OutProf          : 0                        0
Ingress Meter 9 (Multipoint)
For. InProf           : 0                        0
For. OutProf          : 0                        0
===============================================================================
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:7210SAS>show>router#
```

Detailed IP Interface Output — The following table describes the detailed output fields for an IP interface.

Label	Description
If Name	The IP interface name.
Admin State	Down — The IP interface is administratively disabled. Up — The IP interface is administratively enabled.

Label	Description (Continued)
Oper State	Down — The IP interface is operationally disabled. Up — The IP interface is operationally disabled.
IP Addr/mask	The IP address and subnet mask length of the IP interface. Not Assigned — Indicates no IP address has been assigned to the IP interface.
Address Type	Primary — The IP address for the IP interface is the Primary address on the IP interface. Secondary — The IP address for the IP interface is a Secondary address on the IP interface.
IGP Inhibit	Disabled — The secondary IP address on the interface will be recognized as a local interface by the IGP. Enabled — The secondary IP address on the interface will not be recognized as a local interface by the IGP.
Broadcast Address	All-ones — The broadcast format on the IP interface is all ones. Host-ones — The broadcast format on the IP interface is host ones.
If Index	The interface index of the IP router interface.
If Type	Network — The IP interface is a network/core IP interface. Service — The IP interface is a service IP interface.
Port Id	The port ID of the IP interface.
Egress Filter	The egress IP filter policy ID associated with the IP interface. none — Indicates no egress filter policy is associated with the interface.
Ingress Filter	The ingress IP filter policy ID associated with the IP interface. none — Indicates no ingress filter policy is associated with the interface.
QoS Policy	The QoS policy ID associated with the IP interface.
SNTP Broadcast	False — Receipt of SNTP broadcasts on the IP interface is disabled. True — Receipt of SNTP broadcasts on the IP interface is enabled.
MAC Address	The MAC address of the IP interface.
Arp Timeout	The ARP timeout for the interface, in seconds, which is the time an ARP entry is maintained in the ARP cache without being refreshed.
ICMP Mask Reply	False — The IP interface will not reply to a received ICMP mask request. True — The IP interface will reply to a received ICMP mask request.

Label	Description (Continued)
Redirects	Specifies the maximum number of ICMP redirect messages the IP interface will issue in a given period of time (Time (seconds)). Disabled – Indicates the IP interface will not generate ICMP redirect messages.
Unreachables	Specifies the maximum number of ICMP destination unreachable messages the IP interface will issue in a given period of time. Disabled – Indicates the IP interface will not generate ICMP destination unreachable messages.
TTL Expired	The maximum number (Number) of ICMP TTL expired messages the IP interface will issue in a given period of time (Time (seconds)). Disabled – Indicates the IP interface will not generate ICMP TTL expired messages.

```
*A:ALA-12# show router 3 interface detail
===============================================================================
Interface Table
===============================================================================
Interface
-------------------------------------------------------------------------------
If Name      : to-ser1
Admin State  : Up                              Oper State       : Up

IP Addr/mask : 10.10.13.3/24                   Address Type     : Primary
IGP Inhibit  : Disabled                        Broadcast Address: Host-ones

IP Addr/mask : 10.200.0.1/16                   Address Type     : Secondary
IGP Inhibit  : Enabled                         Broadcast Address: Host-ones
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
If Index     : 2
Port Id     : 1/1/2                            If Type          : Network
Egress Filter: none                            Ingress Filter   : 100
QoS Policy   : 1                               SNTP Broadcast   : False
MAC Address  : 04:5d:01:01:00:02               Arp Timeout      : 14400

ICMP Details
Redirects    : Disabled
Unreachables : Number - 100                   Time (seconds)   - 10
TTL Expired  : Number - 100                   Time (seconds)   - 10
===============================================================================
*A:ALA-12#
```

Summary IP Interface Output — The following table describes the summary output fields for the router IP interfaces.

Label	Description
Instance	The router instance number.
Router Name	The name of the router instance.
Interfaces	The number of IP interfaces in the router instance.

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#
```

route-table

Syntax **route-table** [*ip-prefix* [*/mask*] [**longer**] | [**protocol** *protocol*] | [**summary**]]

Context show>router

Description This command displays the active routes in the routing table.

If no command line arguments are specified, all routes are displayed, sorted by prefix.

Parameters *ip-prefix*[/*mask*] **—** Displays routes only matching the specified *ip-prefix* and optional *mask.*

longer — Displays routes matching the *ip-prefix*/*mask* and routes with longer masks.

protocol *protocol* **—** Displays routes learned from the specified protocol.

Values bgp, isis, local, ospf, rip, static, aggregate

summary — Displays a route table summary information.

Output **Standard Show Route Table Output —** The following table describes the standard output fields for the route table.

Label	Description
Dest Address	The route destination address and mask.
Next Hop	The next hop IP address for the route destination.
Type	Local – The route is a local route.
	Remote – The route is a remote route.
Protocol	The protocol through which the route was learned.
Age	The route age in seconds for the route.
Metric	The route metric value for the route.
Pref	The route preference value for the route.
No. of Routes:	The number of routes displayed in the list.

Sample Output

```
*A:ALA-12# show router 3 route-table
===============================================================================
Route Table
===============================================================================
Dest Address       Next Hop        Type    Protocol    Age       Metric  Pref
-------------------------------------------------------------------------------
10.10.0.1/32       10.10.13.1      Remote  OSPF        65844     1001    10
10.10.0.2/32       10.10.13.1      Remote  OSPF        65844     2001    10
10.10.0.3/32       0.0.0.0         Local   Local       1329261   0       0
10.10.0.4/32       10.10.34.4      Remote  OSPF        3523      1001    10
10.10.0.5/32       10.10.35.5      Remote  OSPF        1084022   1001    10
10.10.12.0/24      10.10.13.1      Remote  OSPF        65844     2000    10
10.10.13.0/24      0.0.0.0         Local   Local       65859     0       0
10.10.15.0/24      10.10.13.1      Remote  OSPF        58836     2000    10
10.10.24.0/24      10.10.34.4      Remote  OSPF        3523      2000    10
10.10.25.0/24      10.10.35.5      Remote  OSPF        399059    2000    10
10.10.34.0/24      0.0.0.0         Local   Local       3543      0       0
10.10.35.0/24      0.0.0.0         Local   Local       1329259   0       0
10.10.45.0/24      10.10.34.4      Remote  OSPF        3523      2000    10
10.200.0.0/16      0.0.0.0         Local   Local       4513      0       0
192.168.0.0/20     0.0.0.0         Local   Local       1329264   0       0
192.168.254.0/24   0.0.0.0         Remote  Static      11        1       5
-------------------------------------------------------------------------------
*A:ALA-12#


*A:ALA-12# show router 3 route-table 10.10.0.4
===============================================================================
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#
```

static-arp

Syntax **static-arp** [*ip-address* | *ip-int-name* | **mac** *ieee-mac-addr*]

Context show>router

Description This command displays the router static ARP table sorted by IP address.

If no options are present, all ARP entries are displayed.

Parameters *ip-address* — Only displays static ARP entries associated with the specified IP address.

ip-int-name — Only displays static ARP entries associated with the specified IP interface name.

mac *ieee-mac-addr* **—** Only displays static ARP entries associated with the specified MAC address.

Output **Static ARP Table Output —** The following table describes the output fields for the ARP table.

Label	Description
IP Address	The IP address of the static ARP entry.
MAC Address	The MAC address of the static ARP entry.
Age	The age of the ARP entry. Static ARPs always have 00:00:00 for the age.
Type	Inv – The ARP entry is an inactive static ARP entry (invalid).
	Sta – The ARP entry is an active static ARP entry.
Interface	The IP interface name associated with the ARP entry.
No. of ARP Entries	The number of ARP entries displayed in the list.

Sample Output

```
*A:ALA-12# show router 3 static-arp
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address       Age      Type Interface
-------------------------------------------------------------------------------
10.200.0.253    00:00:5a:40:00:01 00:00:00 Sta  to-ser1
12.200.1.1      00:00:5a:01:00:33 00:00:00 Inv  to-ser1a
-------------------------------------------------------------------------------
No. of ARP Entries: 2
===============================================================================
*A:ALA-12#


*A:ALA-12# show router 3 static-arp 12.200.1.1
===============================================================================
ARP Table
===============================================================================
IP Address    MAC Address        Age     Type Interface
-------------------------------------------------------------------------------
12.200.1.1    00:00:5a:01:00:33 00:00:00 Inv  to-ser1 a
```

```
===============================================================================
*A:ALA-12#

*A:ALA-12# show router 3 static-arp to-ser1
===============================================================================
ARP Table
===============================================================================
IP Address    MAC Address       Age      Type Interface
-------------------------------------------------------------------------------
10.200.0.253  00:00:5a:40:00:01 00:00:00 Sta  to-ser1
===============================================================================
S*A:ALA-12#

*A:ALA-12# show router 3 static-arp mac 00:00:5a:40:00:01
===============================================================================
ARP Table
===============================================================================
IP Address    MAC Address       Age      Type Interface
-------------------------------------------------------------------------------
10.200.0.253  00:00:5a:40:00:01 00:00:00 Sta  to-ser1
===============================================================================
*A:ALA-12#
```

static-route

Syntax **static-route** [*ip-prefix Imask*] | [**preference** *preference*] | [**next-hop** *ip-addr*] [**detail**]

Context show>router

Description This command displays the static entries in the routing table.

If no options are present. all static routes are displayed sorted by prefix.

Parameters *ip-prefix /mask* **—** Displays static routes only matching the specified *ip-prefix* and *mask*.

preference *preference* **—** Only displays static routes with the specified route preference.

Values 0 — 65535

next-hop *ip-addr* **—** Only displays static routes with the specified next hop IP address.

detail — Displays detailed information about the static route.

Output **Show Static Route Output —** The following table describes the output fields for the static route table:

Label	Description
IP Addr/mask	The static route destination address and mask.
Pref	The route preference value for the static route.
Metric	The route metric value for the static route.
Type	BH – The static route is a black hole route. The Nexthop for this type of route is black-hole.

Label	Description (Continued)
	ID — The static route is an indirect route, where the nexthop for this type of route is the non-directly connected next hop. NH — The route is a static route with a directly connected next hop. The Nexthop for this type of route is either the next hop IP address or an egress IP interface name.
Next Hop	The next hop for the static route destination.
Interface	The egress IP interface name for the static route. n/a — indicates there is no current egress interface because the static route is inactive or a black hole route.
Active	N — The static route is inactive; for example, the static route is disabled or the next hop IP interface is down. Y — The static route is active.
No. of Routes:	The number of routes displayed in the list.

Sample Output

```
*A:ALA-12# show router 3 static-route
===============================================================================
Route Table
===============================================================================
IP Addr/mask       Pref Metric Type Nexthop              Interface       Active
-------------------------------------------------------------------------------
192.168.250.0/24   5    1      ID   10.200.10.1          to-ser1         Y
192.168.252.0/24   5    1      NH   10.10.0.254          n/a             N
192.168.253.0/24   5    1      NH   to-ser1              n/a             N
192.168.253.0/24   5    1      NH   10.10.0.254          n/a             N
192.168.254.0/24   4    1      BH   black-hole           n/a             Y
===============================================================================
*A:ALA-12#


*A:ALA-12# show router 3 static-route 192.168.250.0/24
===============================================================================
Route Table
===============================================================================
IP Addr/mask       Pref Metric Type Nexthop              Interface       Active
-------------------------------------------------------------------------------
192.168.250.0/24   5    1      ID   10.200.10.1          to-ser1         Y
===============================================================================
*A:ALA-12#


*A:ALA-12# show router 3 static-route preference 4
===============================================================================
Route Table
===============================================================================
IP Addr/mask       Pref Metric Type Nexthop              Interface       Active
-------------------------------------------------------------------------------
192.168.254.0/24   4    1      BH   black-hole           n/a             Y
===============================================================================
```

```
*A:ALA-12#


*A:ALA-12# show router 3 static-route next-hop 10.10.0.254
===============================================================================
Route Table
===============================================================================
IP Addr/mask       Pref Metric Type Nexthop              Interface       Active
-------------------------------------------------------------------------------
192.168.253.0/24   5    1      NH   10.10.0.254          n/a             N
===============================================================================
*A:ALA-12#

*A:Dut-B# show router static-route

===============================================================================
Static Route Table (Router: Base)  Family: IPv4
===============================================================================
Prefix                                    Tag         Met    Pref Type Act
   Next Hop                                Interface
-------------------------------------------------------------------------------
1.2.3.4/32                                0           1      5    NH   Y
   10.11.25.6
ip-10.11.25.5_base_to_cpe_static
10.11.15.0/24                             0           1      5    NH   Y
   10.11.25.6
ip-10.11.25.5_base_to_cpe_static
-------------------------------------------------------------------------------
No. of Static Routes: 2
===============================================================================


*A:Dut-B# show router static-route detail
===============================================================================
Static Route Table (Router: Base)  Family: IPv4
===============================================================================
Network          : 1.2.3.4/32
Nexthop          : 10.11.25.6
Type             : Nexthop                          Nexthop Type      : IP
Interface        : ip-10.11.25.5_base_to_cpe_stat*  Active            : Y
Metric           : 1                                Preference        : 5
Admin  State     : Up                               Tag               : 0
BFD              : disabled
CPE-check        : enabled                          State             : n/a
Target           : 10.11.18.6
Interval         : 1                                Drop Count        : 3
Log              : N
CPE Host Up Time : 0d 00:00:02
CPE Echo Req Tx  : 3                                CPE Echo Reply Rx : 3
CPE Up Trans     : 1                                CPE Down Trans    : 0
CPE TTL          : 2
-------------------------------------------------------------------------------
Network          : 10.11.15.0/24
Nexthop          : 10.11.25.6
Type             : Nexthop                          Nexthop Type      : IP
Interface        : ip-10.11.25.5_base_to_cpe_stat*  Active            : Y
Metric           : 1                                Preference        : 5
Admin  State     : Up                               Tag               : 0
BFD              : disabled
CPE-check        : disabled
-------------------------------------------------------------------------------
No. of Static Routes: 2
```

```
===============================================================================
```

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.

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.
Metric	The route metric value for the route.

Sample Output

```
*A:ALA-12>config>service#  show router 3 tunnel-table summary
===============================================================================
Tunnel Table Summary (Router: Base)
===============================================================================
                              Active                  Available
-------------------------------------------------------------------------------
LDP                           1                       1
SDP                           1                       1
===============================================================================
*A:ALA-12>config>service#
```

VPRN Clear Commands

arp-host

Syntax **arp-host**
arp-host { **mac** *ieee-address* | **sap** *sap-id* | **ip-address** *ip-address*[*/mask*] }
arp-host [**port** *port-id*] [**inter-dest-id** *intermediate-destination-id* | **no-inter-dest-id**]
arp-host statistics [**sap** *sap-id* | **interface** *interface-name*]

Context clear>service>id

Description This command clears ARP host data.

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 (depending on chassis model)

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-prefixImask*] [**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.

id

Syntax **id** *service-id*

Context clear>service
clear>service>statistics

Description This command clears commands for a specific service.

Parameters *service-id* — The ID that uniquely identifies a service.

Values 1 — 2147483648

sap

Syntax **sap** *sap-id* {**all** | **counters** | **stp**}

Context clear>service>statistics

Description Clears SAP statistics for a SAP.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 987 for command syntax.

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.

VPRN Debug Commands

id

Syntax [**no**] **id** *service-id*

Context debug>service

Description This command debugs commands for a specific service.

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

Parameters *service-id* — The ID that uniquely identifies a service.

sap

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

Context debug>service>id

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

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 987 for command syntax.

sap

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

Context debug>service>id
debug>service>stp

Description This command enables STP debugging for a specific SAP.

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

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 987 for command syntax.

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>id

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>stp

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.

VLL Show Commands

sap-using

Syntax **sap-using** [**sap** *sap-id*]
sap-using interface [*ip-address* | *ip-int-name*]
sap-using [**ingress** | **egress**] **filter** *filter-id*
sap-using [**ingress** | **egress**] **qos-policy** *qos-policy-id*
sap-using encap-type *encap-type*

Context show>service

Description This command displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

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

Values 1.0.0.0 to 223.255.255.255

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

ingress — Specifies matching an ingress policy.

ingress — Specifies matching an ingress policy.

ingress — Specifies matching an ingress policy.

egress — Specifies matching an egress policy.

qos-policy *qos-policy-id* **—** The ingress QoS Policy ID for which to display matching SAPs.

Values 1 — 65535

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. See Common CLI Command Descriptions on page 987 for command syntax.

encap-type *encap-type* — Displays the CEM encapsulation type.

Values cem

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.
MTU	The port MTU value.

Label	Description (Continued)
Ing. QoS	The SAP ingress QoS policy number specified on the ingress SAP.
Ing Fltr	The MAC or IP filter policy ID applied to the ingress SAP.
Egr. QoS	The SAP egress QoS policy number specified on the egress SAP.
Egr. Fltr	The MAC or IP filter policy ID applied to the egress SAP.
Adm	The administrative state of the SAP.
Opr	The operational state of the SAP.

Sample Output

```
*A:Dut-A# show service sap-using

===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.   Ing.    Egr.   Egr.    Adm  Opr
                                           QoS    Fltr    QoS    Fltr
-------------------------------------------------------------------------------
1/1/1:1                         1          1      none    1      none    Up   Up
2/1/2:10/11                     1          1      none    1      none    Up   Up
2/1/2:10/12                     1          1      none    1      none    Up   Up
2/1/2:20/11                     1          1      none    1      none    Up   Up
2/1/2:20/12                     1          1      none    1      none    Up   Up
2/1/4:cp.10                     10         1      none    1      none    Up   Up
2/1/4:cp.20                     20         1      none    1      none    Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 7
-------------------------------------------------------------------------------
===============================================================================
A:Dut-A>config>service>vpls# show service sap-using
===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.   Ing.    Egr.    Adm  Opr
                                           QoS    Fltr    Fltr
-------------------------------------------------------------------------------
lag-3:100                       100        1      none    none    Up   Up
1/1/3                           101        10     mac     none    Up   Up
lag-3:101                       101        10     mac     none    Up   Up
lag-3:102                       102        10     mac     none    Up   Up
lag-3:103                       103        10     mac     none    Up   Up
lag-3:104                       104        10     mac     none    Up   Up
lag-3:105                       105        10     mac     none    Up   Up
lag-3:201                       201        10     mac     none    Up   Up
lag-3:202                       202        10     mac     none    Up   Up
lag-3:203                       203        10     mac     none    Up   Up
lag-3:204                       204        10     mac     none    Up   Up
lag-3:205                       205        10     mac     none    Up   Up
1/1/16:301                      301        10     mac     none    Up   Up
lag-4:301                       301        10     mac     none    Up   Up
1/1/16:302                      302        10     mac     none    Up   Up
```

```
lag-4:302                       302        10    mac     none   Up   Up
1/1/16:303                      303        10    mac     none   Up   Up
lag-4:303                       303        10    mac     none   Up   Up
1/1/16:304                      304        10    mac     none   Up   Up
lag-4:304                       304        10    mac     none   Up   Up
1/1/16:305                      305        10    mac     none   Up   Up
lag-4:305                       305        10    mac     none   Up   Up
...
===============================================================================
A:Dut-A>config>service>vpls#


A:Dut-A>config>service# show service sap-using sap 1/1/16:305
===============================================================================
Service Access Points Using Port 1/1/16:305
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.   Adm  Opr
                                           QoS   Fltr    Fltr
-------------------------------------------------------------------------------
1/1/16:305                      305        10    mac     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 1
===============================================================================
A:Dut-A>config>service#


A:ces-A# show service sap-using sap 1/2/1.1
===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.   Adm  Opr
                                           QoS   Fltr    Fltr
-------------------------------------------------------------------------------
1/2/1.1                         1          12    none    none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
A:ces-A#

*A:ces-A# show service sap-using sap 1/2/1.1

===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.   Adm  Opr
                                           QoS   Fltr    Fltr
-------------------------------------------------------------------------------
1/2/1.1                         1          1     none    none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================


*A:ces-A# show service sap-using encap-type cem

===============================================================================
Service Access Points Using Encap Type 'cem'
===============================================================================
PortId                            SvcId           Adm      Opr      Alarm
```

```
-------------------------------------------------------------------------------
1/2/1.1                          1              Up       Up       No
1/2/2.1                          2              Up       Up       No
1/2/3.1                          3              Up       Down     Yes
1/2/4.1                          4              Up       Down     Yes
-------------------------------------------------------------------------------
Number of SAPS : 4
-------------------------------------------------------------------------------
===============================================================================
```

sdp

Note : SDP commands are not supported by 7210 SAS-M devices configured in uplink mode.

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

Context show>service

Description This command displays SDP information.

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

Parameters *sdp-id* — The SDP ID for which to display information.

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 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 Delivery	Specifies the type of delivery used by the SDP: MPLS.
Flags	Specifies all the conditions that affect the operating status of this SDP.

Label	Description (Continued)
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.
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.

Sample Output

```
*A:ALA-12# show service sdp
===============================================================================
Services: Service Destination Points
===============================================================================
SdpId    Adm MTU   Opr MTU   IP address       Adm  Opr         Deliver Signal
-------------------------------------------------------------------------------
10       4462      4462      10.20.1.3        Up   Dn NotReady MPLS    TLDP
40       4462      1534      10.20.1.20       Up   Up          MPLS    TLDP
60       4462      1514      10.20.1.21       Up   Up          MPLS TLDP
100      4462      4462      180.0.0.2        Down Down        MPLS TLDP
500      4462      4462      10.20.1.50       Up   Dn NotReady MPLS TLDP
-------------------------------------------------------------------------------
Number of SDPs : 5
===============================================================================
*A:ALA-12#


*A:ALA-12# show service sdp 2 detail
===============================================================================
Service Destination Point (Sdp Id : 2) Details
===============================================================================
Sdp Id 2  -(10.10.10.104)
-------------------------------------------------------------------------------
Description           : MPLS-10.10.10.104
SDP Id               : 2
Admin Path MTU       : 0                    Oper Path MTU      : 0
Far End              : 10.10.10.104         Delivery           : MPLS
Admin State          : Up                   Oper State         : Down
Flags                : SignalingSessDown TransportTunnDown
Signaling            : TLDP                 VLAN VC Etype      : 0x8100
Last Status Change   : 02/01/2007 09:11:39  Adv. MTU Over.     : No
Last Mgmt Change     : 02/01/2007 09:11:46

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

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
===============================================================================
*A:ALA-12#
*A:ALA-12# show service sdp 8
===============================================================================
Service Destination Point (Sdp Id : 8)
===============================================================================
SdpId    Adm MTU   Opr MTU   IP address       Adm  Opr         Deliver Signal
-------------------------------------------------------------------------------
8        4462      4462      10.10.10.104     Up   Dn NotReady MPLS    TLDP
===============================================================================
*A:ALA-12#


*A:ALA-12# show service sdp 8 detail
===============================================================================
Service Destination Point (Sdp Id : 8) Details
===============================================================================
Sdp Id 8  -(10.10.10.104)
-------------------------------------------------------------------------------
```

```
Description           : MPLS-10.10.10.104
SDP Id               : 8
Admin Path MTU       : 0                    Oper Path MTU      : 0
Far End              : 10.10.10.104         Delivery           : MPLS
Admin State          : Up                   Oper State         : Down
Flags                : SignalingSessDown TransportTunnDown
Signaling            : TLDP                 VLAN VC Etype      : 0x8100
Last Status Change   : 02/01/2007 09:11:39  Adv. MTU Over.     : No
Last Mgmt Change     : 02/01/2007 09:11:46

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

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

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 —** 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:ces-A# show service sdp-using
===============================================================================
SDP Using
===============================================================================
SvcId       SdpId               Type    Far End        Opr S* I.Label  E.Label
-------------------------------------------------------------------------------
1           12:1                Spok    2.2.2.2        Up     131063   131062
2           12:2                Spok    2.2.2.2        Up     131062   131069
3           122:3               Spok    2.2.2.2        Up     131069   131068
4           12:4                Spok    2.2.2.2        Up     131061   131061
-------------------------------------------------------------------------------
Number of SDPs : 4
-------------------------------------------------------------------------------
===============================================================================
*A:ces-A#
```

service-using

Syntax **service-using** [**cpipe**] [**sdp** *sdp-id*] [**b-vpls**] [**i-vpls**] [**m-vpls**] [**sdp** *sdp-id*] [**customer** *customer-id*]

Context show>service

Description This command displays the services matching certain usage properties.

If no optional parameters are specified, all services defined on the system are displayed.

Parameters [**service**] — Displays information for the specified service type.

b-vpls — Specifies the B-component instance of the Provider Backbone Bridging (PBB/IEEE 802.1ah) feature. It represents the multi-point tunneling component that multiplexes multiple customer VPNs (ISIDs) together. It is similar to a regular VPLS instance that operates on the backbone MAC addresses.

i-vpls — Specifies the I-component instance of the Provider Backbone Bridging (PBB/IEEE 802.1ah) feature. It identifies the specific VPN entity associated to a customer multipoint (ELAN) service. It is similar to a regular VPLS instance that operates on the customer MAC addresses.

m-vpls — Specifies the M-component (managed VPLS) instance of the Provider Backbone Bridging (PBB/IEEE 802.1ah) feature.

sdp *sdp-id* **—** Displays only services bound to the specified SDP ID.

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-using output —** The following table describes the command 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
300          Epipe     Up     Up         10                09/05/2006 13:24:15
-------------------------------------------------------------------------------
Matching Services : 2
===============================================================================
*A:ALA-12#


*A:ALA-12# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
1            uVPLS     Up     Up         1                 10/26/2006 15:44:57
2            Epipe     Up     Down       1                 10/26/2006 15:44:57
10           mVPLS     Down   Down       1                 10/26/2006 15:44:57
11           mVPLS     Down   Down       1                 10/26/2006 15:44:57
100          mVPLS     Up     Up         1                 10/26/2006 15:44:57
101          mVPLS     Up     Up         1                 10/26/2006 15:44:57
102          mVPLS     Up     Up         1                 10/26/2006 15:44:57
999          uVPLS     Down   Down       1                 10/26/2006 16:14:33
-------------------------------------------------------------------------------
Matching Services : 8
-------------------------------------------------------------------------------
*A:ALA-12#
*A:ces-A# show service service-using cpipe

===============================================================================
Services [cpipe]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
1            Cpipe     Up     Up         1                 05/20/2010 00:12:16
2            Cpipe     Up     Up         1                 05/20/2010 00:12:17
3            Cpipe     Up     Down       1                 05/20/2010 00:12:17
4            Cpipe     Up     Down       1                 05/20/2010 00:12:17
-------------------------------------------------------------------------------
Matching Services : 4
-------------------------------------------------------------------------------
===============================================================================
```

id

Syntax **id** *service-id* {**all** | **arp** | **base** | **endpoint** | **fdb** | **interface** | **label** | **labels** | **sap** | **split-horizon-group** | **stp| interface | mstp-configuration**}

Context show>service

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

Parameters *service-id* — The service identification number that identifies the service in the domain.

Values service-id: 1 — 214748364
svc-name: A string up to 64 characters in length.

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 information.

interface — Display service interfaces.

labels — Display labels being used by this service.

mstp-configuration — Display MSTP information.

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.

Sample Output

```
*A:ces-A# show service id 1 sap
===============================================================================
SAP(Summary), Service 1
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.   Adm  Opr
                                           QoS   Fltr    Fltr
-------------------------------------------------------------------------------
1/2/1.1                         1          1     none    none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================

*A:ces-A# show service id 1 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                  Vpn Id           : 0
Service Type      : Cpipe              VLL Type         : SAToPT1
Description       : (Not Specified)
Customer Id       : 1
```

```
Last Status Change: 07/06/2010 19:21:14
Last Mgmt Change  : 07/06/2010 19:21:14
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/2/1.1                              cem          1514    1514    Up   Up
sdp:12:1 S(2.2.2.2)                      n/a          0       9190    Up   Up
===============================================================================
*A:Dut-A>show# service id 104 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 104                 Vpn Id            : 0
Service Type      : Cpipe               VLL Type          : CESoPSN
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 12/15/2010 07:39:05
Last Mgmt Change  : 12/15/2010 07:25:37
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/2/1.2                              cem          1514    1514    Up   Up
sdp:123:104 S(102.102.102.102)           n/a          0       9190    Up   Up
===============================================================================
*A:Dut-A>show# service id 104 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 104                 Vpn Id            : 0
Service Type      : Cpipe               VLL Type          : CESoPSN
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 12/15/2010 07:39:05
Last Mgmt Change  : 12/15/2010 07:25:37
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/2/1.2                              cem          1514    1514    Up   Up
sdp:123:104 S(102.102.102.102)           n/a          0       9190    Up   Up
```

```
===============================================================================
*A:Dut-A>show#
```

all

Syntax **all**

Context show>service>id

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

Output **Show service ID Output —** The following table describes the output fields when the **all** option is specified:

Label	Description
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.
Endpoint	Specifies the name of the service endpoint.
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, L2OperDown, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
SAP Count	The number of SAPs specified for this service.
SDP Bind Count	The number of SDPs bound to this service.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.
Type	Indicates whether this Service SDP binding is a spoke or a mesh.
Admin Path MTU	The 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 Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Jitter Buffer (packets)	Indicates the jitter buffer length in number of packet buffers.
Playout Threshold (packets)	Indicates the playout buffer packets threshold in number of packet buffers.
Playout Threshold (packets)	Indicates the current packet depth of the jitter buffer.
Peer Pw Bits	Indicates the bits set by the LDP peer when there is a fault on its side of the pseudowire. LAC failures occur on the SAP that has been configured on the pipe service, PSN bits are set by SDP-binding failures on the pipe service. The pwNotForwarding bit is set when none of the above failures apply, such as an MTU mismatch failure. This value is only applicable if the peer is using the pseudowire status signalling method to indicate faults. pwNotForwarding — Pseudowire not forwarding lacIngressFault Local — Attachment circuit RX fault lacEgresssFault Local — Attachment circuit TX fault psnIngressFault Local — PSN-facing PW RX fault psnEgressFault Local — PSN-facing PW TX fault pwFwdingStandby — Pseudowire in standby mode
Signaling Override	Indicates the overriding signaled pseudowire type, as configured under the **signaled-vc-type-override** option for Apipes. This field is only displayed if **signaled-vc-type-override** is configured.
LLF Admin State	Displays the Link Loss Forwarding administrative state.
LLF Oper State	Displays the Link Loss Forwarding operational state.
Standby Signaling Master	Indicates if the parameter standby signalling master is enabled.

Sample Output

```
*A:Dut-A>show>service>id# all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 1501               Vpn Id            : 1501
Service Type      : Epipe
Description       : Default epipe description for service id 1501
```

```
Customer Id       : 1
Last Status Change: 02/21/2011 13:07:03
Last Mgmt Change  : 02/21/2011 13:03:58
Admin State       : Up                  Oper State        : Up
MTU               : 1514
MTU Check         : Enabled
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 2
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 1413:1501  -(10.20.1.4)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id            : 1413:1501               Type              : Spoke
VC Type           : Ether                   VC Tag            : n/a
Admin Path MTU    : 0                       Oper Path MTU     : 9182
Far End           : 10.20.1.4               Delivery          : MPLS

Admin State       : Up                      Oper State        : Up
Acct. Pol         : 14                      Collect Stats     : Enabled
Ingress Label     : 130948                  Egress Label      : 130483
Ing mac Fltr      : n/a                     Egr mac Fltr      : n/a
Ing ip Fltr       : n/a                     Egr ip Fltr       : n/a
Admin ControlWord : Preferred               Oper ControlWord  : True
Admin BW(Kbps)    : 0                       Oper BW(Kbps)     : 0
Last Status Change : 02/21/2011 13:07:12    Signaling         : TLDP
Last Mgmt Change  : 02/21/2011 13:03:58     Force Vlan-Vc     : Disabled
Endpoint          : coreSide                Precedence        : 1
Class Fwding State : Down
Flags             : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : pwe3ControlWord mplsRouterAlertLabel

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

Statistics          :
I. Fwd. Pkts.     : 48319                   I. Fwd. Octs.     : 5690869
E. Fwd. Pkts.     : 34747                   E. Fwd. Octets    : 4013709
-------------------------------------------------------------------------------
Eth-Cfm Configuration Information
-------------------------------------------------------------------------------
Md-index          : 1000                    Direction         : Down
Ma-index          : 1150114                 Admin             : Enabled
MepId             : 1                       CCM-Enable        : Enabled
LowestDefectPri   : macRemErrXcon           HighestDefect     : none
Defect Flags      : None
Mac Address       : 7c:20:64:ad:04:07       ControlMep        : False
CcmLtmPriority    : 7
CcmTx             : 11385                   CcmSequenceErr    : 0
Eth-1Dm Threshold : 3(sec)
Eth-Ais:          : Disabled
Eth-Tst:          : Disabled
LbRxReply         : 0                       LbRxBadOrder      : 0
LbRxBadMsdu       : 0                       LbTxReply         : 0
```

```
LbNextSequence     : 1                        LtNextSequence    : 1
LtRxUnexplained    : 0

Associated LSP LIST :
Lsp Name           : A_D_21
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 03h49m30s

-------------------------------------------------------------------------------
 Sdp Id 1613:1501  -(10.20.1.6)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 1613:1501                Type              : Spoke
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 9182
Far End            : 10.20.1.6                Delivery          : MPLS

Admin State        : Up                       Oper State        : Up
Acct. Pol          : 14                       Collect Stats     : Enabled
Ingress Label      : 130526                   Egress Label      : 130424
Ing mac Fltr       : n/a                      Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                      Egr ip Fltr       : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Admin BW(Kbps)     : 0                        Oper BW(Kbps)     : 0
Last Status Change : 02/21/2011 13:07:03      Signaling         : TLDP
Last Mgmt Change   : 02/21/2011 13:03:58      Force Vlan-Vc     : Disabled
Endpoint           : coreSide                 Precedence        : 2
Class Fwding State : Down
Flags              : None
Peer Pw Bits       : pwFwdingStandby
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : mplsRouterAlertLabel

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

Statistics           :
I. Fwd. Pkts.      : 25                       I. Fwd. Octs.     : 2776
E. Fwd. Pkts.      : 23                       E. Fwd. Octets    : 2557
-------------------------------------------------------------------------------
Eth-Cfm Configuration Information
-------------------------------------------------------------------------------
Md-index           : 1000                     Direction         : Down
Ma-index           : 1150116                  Admin             : Enabled
MepId              : 1                        CCM-Enable        : Enabled
LowestDefectPri    : macRemErrXcon            HighestDefect     : none
Defect Flags       : None
Mac Address        : 7c:20:64:ad:04:07        ControlMep        : False
CcmLtmPriority     : 7
CcmTx              : 11414                    CcmSequenceErr    : 0
Eth-1Dm Threshold  : 3(sec)
Eth-Ais:           : Disabled
Eth-Tst:           : Disabled
LbRxReply          : 0                        LbRxBadOrder      : 0
LbRxBadMsdu        : 0                        LbTxReply         : 0
LbNextSequence     : 1                        LtNextSequence    : 1
LtRxUnexplained    : 0
```

```
Associated LSP LIST :
Lsp Name           : A_F_21
Admin State        : Up                        Oper State        : Up
Time Since Last Tr*: 03h48m45s

-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP lag-3:1501.1501
-------------------------------------------------------------------------------
Service Id         : 1501
SAP                : lag-3:1501.1501          Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                        Oper State        : Up
Flags              : None
Last Status Change : 02/21/2011 13:06:45
Last Mgmt Change   : 02/21/2011 13:03:58

Admin MTU          : 9212                      Oper MTU          : 9212
Ingr IP Fltr-Id    : n/a                       Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : 1501                      Egr Mac Fltr-Id   : n/a
tod-suite          : None
Egr Agg Rate Limit : max
Endpoint           : accessSide

Acct. Pol          : Default                   Collect Stats     : Enabled


-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1500                      Egress qos-policy : 1500
-------------------------------------------------------------------------------
Sap Egress Policy (1500)
-------------------------------------------------------------------------------
Scope              : Template
Remark             : False                     Remark Pol Id     : 2
Accounting         : frame-based
Description    : Sap Egress Policy for svcList 1500
-------------------------------------------------------------------------------
Queue Rates and Rules
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
QueueId        CIR          CIR Adpt Rule        PIR          PIR Adpt Rule
-------------------------------------------------------------------------------
Queue1          10000          max                10000          max
Queue2          10000          max                10000          max
Queue3          10000          max                10000          max
Queue4          10000          max                10000          max
Queue5          10000          max                10000          max
Queue6          10000          max                10000          max
Queue7          10000          max                10000          max
Queue8          10000          max                10000          max
-------------------------------------------------------------------------------
Parent Details
```

```
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
QueueId          Port       CIR Level      PIR Weight
-------------------------------------------------------------------------------
Queue1          True             1             1
Queue2          True             2             2
Queue3          True             3             3
Queue4          True             4             4
Queue5          True             5             5
Queue6          True             6             6
Queue7          True             7             7
Queue8          True             8             8
-------------------------------------------------------------------------------
High Slope
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
QueueId        State      Start-Avg(%)   Max-Avg(%)   Max-Prob(%)
-------------------------------------------------------------------------------
Queue1           Up             50           100            50
Queue2           Up             50           100            50
Queue3           Up             50           100            50
Queue4           Up             50           100            50
Queue5           Up             50           100            50
Queue6           Up             50           100            50
Queue7           Up             50           100            50
Queue8           Up             50           100            50
-------------------------------------------------------------------------------
Low Slope
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
QueueId        State      Start-Avg(%)   Max-Avg(%)   Max-Prob(%)
-------------------------------------------------------------------------------
Queue1           Up             10           50            50
Queue2           Up             10           50            50
Queue3           Up             10           50            50
Queue4           Up             10           50            50
Queue5           Up             10           50            50
Queue6           Up             10           50            50
Queue7           Up             10           50            50
Queue8           Up             10           50            50
-------------------------------------------------------------------------------
Burst Sizes and Time Average Factor
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
QueueId         CBS            MBS       Time Average Factor        Queue-Mgmt
-------------------------------------------------------------------------------
Queue1          200              400             10           qM_1500
Queue2          200              400             10           qM_1500
Queue3          200              400             10           qM_1500
Queue4          200              400             10           qM_1500
Queue5          200              400             10           qM_1500
Queue6          200              400             10           qM_1500
Queue7          200              400             10           qM_1500
Queue8          200              400             10           qM_1500
-------------------------------------------------------------------------------
Aggregate Policer (Available)
-------------------------------------------------------------------------------
rate               : n/a                      burst            : n/a
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
```

```
Classifiers Allocated: 32                   Meters Allocated   : 16
Classifiers Used     : 8                    Meters Used        : 5
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress Stats:          34659                   3241035
Egress Stats:           48099                   5291928
Extra-Tag Drop Stats:   n/a                     n/a
-------------------------------------------------------------------------------
Sap per Meter stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Meter 1 (Unicast)
For. InProf           : 7209                    468585
For. OutProf          : 0                       0

Ingress Meter 2 (Unicast)
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Meter 3 (Unicast)
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Meter 4 (Unicast)
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Meter 5 (Unicast)
For. InProf           : 27454                   2772854
For. OutProf          : 0                       0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Egress Queue 1 (be)
Fwd Stats             :  0                      0
Drop InProf           :  0                      0
Drop OutProf          :  0                      0

Egress Queue 2 (l2)
Fwd Stats             :  3                      180
Drop InProf           :  0                      0
Drop OutProf          :  0                      0

Egress Queue 3 (af)
Fwd Stats             :  0                      0
Drop InProf           :  0                      0
Drop OutProf          :  0                      0

Egress Queue 4 (l1)
Fwd Stats             :  0                      0
Drop InProf           :  0                      0
Drop OutProf          :  0                      0

Egress Queue 5 (h2)
Fwd Stats             :  0                      0
Drop InProf           :  0                      0
```

```
Drop OutProf         :  0                      0

Egress Queue 6 (ef)
Fwd Stats            :  0                      0
Drop InProf          :  0                      0
Drop OutProf         :  0                      0

Egress Queue 7 (h1)
Fwd Stats            :  0                      0
Drop InProf          :  0                      0
Drop OutProf         :  0                      0

Egress Queue 8 (nc)
Fwd Stats            :  20842                  1938306
Drop InProf          :  0                      0
Drop OutProf         :  0                      0

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
Endpoint name               : coreSide
Description                 : (Not Specified)
Revert time                 : 0
Act Hold Delay              : 0
Standby Signaling Master    : true
Tx Active                   : 1413:1501
Tx Active Up Time           : 0d 03:48:41
Revert Time Count Down      : N/A
Tx Active Change Count      : 2
Last Tx Active Change       : 02/21/2011 13:07:12
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
Spoke-sdp: 1413:1501 Prec:1                       Oper Status: Up
Spoke-sdp: 1613:1501 Prec:2                       Oper Status: Up
===============================================================================
Endpoint name               : accessSide
Description                 : (Not Specified)
Revert time                 : 0
Act Hold Delay              : 0
Standby Signaling Master    : false
Tx Active                   : lag-3:1501.1501
Tx Active Up Time           : 0d 03:49:08
Revert Time Count Down      : N/A
Tx Active Change Count      : 1
Last Tx Active Change       : 02/21/2011 13:06:45
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
SAP      : lag-3:1501.1501                        Oper Status: Up
===============================================================================
===============================================================================


*A:ces-A# show service id 1 all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 1                  Vpn Id           : 0
Service Type      : Cpipe              VLL Type         : SAToPT1
```

```
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 07/06/2010 19:21:14
Last Mgmt Change  : 07/06/2010 19:21:14
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 12:1  -(2.2.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 12:1                      Type              : Spoke
VC Type           : SAToPT1                   VC Tag            : 0
Admin Path MTU    : 0                         Oper Path MTU     : 9190
Far End           : 2.2.2.2                   Delivery          : MPLS

Admin State       : Up                        Oper State        : Up
Acct. Pol         : None                      Collect Stats     : Disabled
Ingress Label     : 131064                    Egress Label      : 131064
Admin ControlWord : Preferred                 Oper ControlWord  : True
Admin BW(Kbps)    : 0                         Oper BW(Kbps)     : 0
Last Status Change : 07/06/2010 19:21:14      Signaling         : TLDP
Last Mgmt Change  : 07/06/2010 19:21:14
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

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.     : 141578                    I. Fwd. Octs.     : 31430316
E. Fwd. Pkts.     : 141583                    E. Fwd. Octets    : 31431426

Associated LSP LIST :
Lsp Name          : to_b_1_2
Admin State       : Up                        Oper State        : Up
Time Since Last Tr*: 04h08m22s
-------------------------------------------------------------------------------
CPIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Local Bit-rate    : 24                        Peer Bit-rate     : 24
Local Payload Size : 192                      Peer Payload Size : 192
Local Sig Pkts    : No Sig.                   Peer Sig Pkts     : No Sig.
Local CAS Framing : No CAS                    Peer CAS Framing  : No CAS
Local RTP Header  : No                        Peer RTP Header   : No
Local Differential : No                       Peer Differential : No
Local Timestamp   : 0                         Peer Timestamp    : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
```

```
Service Access Points
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
SAP 1/2/1.1
-------------------------------------------------------------------------------
Service Id         : 1
SAP                : 1/2/1.1                  Encap             : cem
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Up
Flags              : None
Last Status Change : 07/06/2010 14:16:41
Last Mgmt Change   : 07/06/2010 11:31:34

Admin MTU          : 1514                     Oper MTU          : 1514
Endpoint           : N/A

Acct. Pol          : None                     Collect Stats     : Disabled
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                       Packets                 Octets
Ingress Stats:         705193                  153732074
Egress Stats:          705179                  153729022
-------------------------------------------------------------------------------
CEM SAP Configuration Information
-------------------------------------------------------------------------------
Endpoint Type          : Unstruct. T1     Bit-rate                   : 24
Payload Size           : 192              Jitter Buffer (ms)         : 5
Jitter Buffer (packets): 6                Playout Threshold (packets): 4
Use RTP Header         : No               Differential               : No
Timestamp Freq         : 0                CAS Framing                : No CAS
Effective PDVT         : +/-2.984 ms

Cfg Alarm      : stray malformed pktloss overrun underrun
Alarm Status   :
-------------------------------------------------------------------------------
CEM SAP Statistics
-------------------------------------------------------------------------------
                         Packets           Seconds           Events
Egress Stats
Forwarded             :  705523
Dropped               :  0
Missing               :  0
Reordered Forwarded   :  0
Underrun              :  11119                                3
Overrun               :  0                                    0
Misordered Dropped    :  0
Malformed Dropped     :  0
LBit Dropped          :  0
Multiple Dropped      :  0
Error                 :                    17
Severely Error        :                    15
Unavailable           :                    0
Failure Count         :                                       1
Jitter Buffer Depth   :  3

Ingress Stats
```

```
Forwarded             :  705574
Dropped               :  0
-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================

===============================================================================
*A:Dut-A>show# service id 104 all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 104                 Vpn Id            : 0
Service Type      : Cpipe               VLL Type          : CESoPSN
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 12/15/2010 07:39:05
Last Mgmt Change  : 12/15/2010 07:25:37
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 123:104  -(102.102.102.102)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id              : 123:104                 Type              : Spoke
VC Type             : CESoPSN                 VC Tag            : 0
Admin Path MTU      : 0                       Oper Path MTU     : 9190
Far End             : 102.102.102.102         Delivery          : MPLS

Admin State         : Up                      Oper State        : Up
Acct. Pol           : None                    Collect Stats     : Disabled
Ingress Label       : 131069                  Egress Label      : 131068
Admin ControlWord   : Preferred               Oper ControlWord  : True
Admin BW(Kbps)      : 0                       Oper BW(Kbps)     : 0
Last Status Change  : 12/15/2010 07:27:17     Signaling         : TLDP
Last Mgmt Change    : 12/15/2010 07:25:37
Endpoint            : y                       Precedence        : 4
Flags               : None
Peer Pw Bits        : None
Peer Fault Ip       : None
Peer Vccv CV Bits   : lspPing
Peer Vccv CC Bits   : pwe3ControlWord mplsRouterAlertLabel

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

Statistics            :
I. Fwd. Pkts.       : 770680                  I. Fwd. Octs.     : 72443920
E. Fwd. Pkts.       : 772901                  E. Fwd. Octets    : 72652694

Associated LSP LIST :
Lsp Name            : static-32
```

```
Admin State        : Up                        Oper State        : Up
Time Since Last Tr*: 01h55m01s


-------------------------------------------------------------------------------
CPIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Local Bit-rate     : 1                         Peer Bit-rate     : 1
Local Payload Size : 64                        Peer Payload Size : 64
Local Sig Pkts     : No Sig.                   Peer Sig Pkts     : No Sig.
Local CAS Framing  : No CAS                    Peer CAS Framing  : No CAS
Local RTP Header   : No                        Peer RTP Header   : No
Local Differential : No                        Peer Differential : No
Local Timestamp    : 0                         Peer Timestamp    : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/2/1.2
-------------------------------------------------------------------------------
Service Id         : 104
SAP                : 1/2/1.2                   Encap             : cem
Description        : (Not Specified)
Admin State        : Up                        Oper State        : Up
Flags              : None
Last Status Change : 12/15/2010 07:39:05
Last Mgmt Change   : 12/15/2010 07:25:37

Admin MTU          : 1514                      Oper MTU          : 1514
Endpoint           : N/A

Acct. Pol          : None                      Collect Stats     : Disabled


-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                         Egress qos-policy : 1
-------------------------------------------------------------------------------
Aggregate Policer
-------------------------------------------------------------------------------
rate               : n/a                       burst             : n/a
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                       Packets                   Octets
Ingress Stats:         773839                    69645510
Egress Stats:          771668                    69450120
Extra-Tag Drop Stats:  n/a                       n/a
-------------------------------------------------------------------------------
CEM SAP Configuration Information
-------------------------------------------------------------------------------
Endpoint Type          : NxDS0           Bit-rate                  : 1
Payload Size           : 64              Jitter Buffer (ms)        : 32
Jitter Buffer (packets): 4               Playout Threshold (packets): 3
Use RTP Header         : No              Differential              : No
Timestamp Freq         : 0               CAS Framing               : No CAS
```

```
Effective PDVT         : +/-16.0 ms

Cfg Alarm      : stray malformed pktloss overrun underrun
Alarm Status   :
-------------------------------------------------------------------------------
CEM SAP Statistics
-------------------------------------------------------------------------------
                          Packets           Seconds           Events
Egress Stats
Forwarded             :  771800
Dropped               :  132
Missing               :  0
Reordered Forwarded   :  0
Underrun              :  2355                                 1
Overrun               :  0                                    0
Misordered Dropped    :  0
Malformed Dropped     :  0
LBit Dropped          :  132
Multiple Dropped      :  0
Error                 :                    1
Severely Error        :                    0
Unavailable           :                    18
Failure Count         :                                       1
Jitter Buffer Depth   :  2

Ingress Stats
Forwarded             :  774156
Dropped               :  0

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
Endpoint name               : y
Description                 : (Not Specified)
Revert time                 : 0
Act Hold Delay              : 0
Tx Active                   : 123:104
Tx Active Up Time           : 0d 01:55:06
Revert Time Count Down      : N/A
Tx Active Change Count      : 1
Last Tx Active Change       : 12/15/2010 07:27:17
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
Spoke-sdp: 123:104 Prec:4                          Oper Status: Up
===============================================================================
===============================================================================
*A:Dut-A>show#
```

Sample output (Meter-override)

```
A:7210SAS>show>service# id 1101 sap 1/2/1:1 detail
Ingress Meter Override
-------------------------------------------------------------------------------

Meter Id          : 1
Admin PIR         : 12000                  Admin CIR        : 10000
Oper PIR          : 12000                  Oper CIR         : 10000
PIR Rule          : closest*               CIR Rule         : closest*
```

```
MBS                : max*                      CBS                : max*
Mode               : Trtcm2*

* means the value is inherited
-------------------------------------------------------------------------------
A:7210SAS>show>service#
```

base

Syntax **base**

Context show>service>id

Description Displays basic information about the service ID including service type, description, SAPs.

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, VPLS
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Adm	The desired state of the service.
Oper	The operating state of the service.
Mtu	The largest frame size (in octets) that the service can handle.
Def. Mesh VC Id	This object is only valid in services that accept mesh SDP bindings. It is used to validate the VC ID portion of each mesh SDP binding defined in the service.
SAP Count	The number of SAPs defined on the service.
SDP Bind Count	The number of SDPs bound to the service.
Identifier	Specifies the service access (SAP) points.
Type	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received.
AdmMTU	Specifies the desired largest service frame size (in octets) that can be transmitted through this SAP, without requiring the packet to be fragmented.
PBB Tunnel Point	Specifies the endpoint in the B-VPLS environment where the Epipe terminates.
Admin MTU	Specifies the B-VPLS admin MTU.
Backbone-Flooding	Specifies whether or not the traffic is flooded in the B-VPLS for the detination instead of unicast. If the backbone destination MAC is in the B-VPLS FDB, then it will be unicast.

Label	Description (Continued)
ISID	The 24 bit field carrying the service instance identifier associated with the frame. It is used at the destination PE as a demultiplexor field.

Sample Output

```
A:Dut-A# show service id 1101 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1101               Vpn Id            : 1101
Service Type      : Epipe
Description       : Default epipe description for service id 1101
Customer Id       : 1
Last Status Change: 07/07/2009 18:13:43
Last Mgmt Change  : 07/07/2009 14:39:14
Admin State       : Up                 Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                  SDP Bind Count    : 1
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:lag-4:1101                         q-tag        9212    9212    Up   Up
sdp:1409:1101 S(10.20.1.4)             n/a          0       9186    Up   Up
===============================================================================
A:Dut-A#
*A:ces-A# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                  Vpn Id            : 0
Service Type      : Cpipe              VLL Type          : SAToPT1
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 07/06/2010 19:21:14
Last Mgmt Change  : 07/06/2010 19:21:14
Admin State       : Up                 Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                  SDP Bind Count    : 1

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/2/1.1                            cem          1514    1514    Up   Up
sdp:12:1 S(2.2.2.2)                    n/a          0       9190    Up   Up
===============================================================================

*A:Dut-A>show# service id 104 base

===============================================================================
```

```
Service Basic Information
===============================================================================
Service Id        : 104                Vpn Id            : 0
Service Type      : Cpipe              VLL Type          : CESoPSN
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 12/15/2010 07:39:05
Last Mgmt Change  : 12/15/2010 07:25:37
Admin State       : Up                 Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                  SDP Bind Count    : 1

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/2/1.2                            cem          1514    1514    Up   Up
sdp:123:104 S(102.102.102.102)         n/a          0       9190    Up   Up
===============================================================================
*A:Dut-A>show# service id 104 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 104                Vpn Id            : 0
Service Type      : Cpipe              VLL Type          : CESoPSN
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 12/15/2010 07:39:05
Last Mgmt Change  : 12/15/2010 07:25:37
Admin State       : Up                 Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                  SDP Bind Count    : 1

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/2/1.2                            cem          1514    1514    Up   Up
sdp:123:104 S(102.102.102.102)         n/a          0       9190    Up   Up
===============================================================================
*A:Dut-A>show#
```

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:Dut-A>show>service>id# endpoint

===============================================================================
Service 1501 endpoints
===============================================================================
Endpoint name               : coreSide
Description                 : (Not Specified)
Revert time                 : 0
Act Hold Delay              : 0
Standby Signaling Master    : true
Tx Active                   : 1413:1501
Tx Active Up Time           : 0d 03:46:25
Revert Time Count Down      : N/A
Tx Active Change Count      : 2
Last Tx Active Change       : 02/21/2011 13:07:12
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
Spoke-sdp: 1413:1501 Prec:1                       Oper Status: Up
Spoke-sdp: 1613:1501 Prec:2                       Oper Status: Up
===============================================================================
Endpoint name               : accessSide
Description                 : (Not Specified)
Revert time                 : 0
Act Hold Delay              : 0
Standby Signaling Master    : false
Tx Active                   : lag-3:1501.1501
Tx Active Up Time           : 0d 03:46:52
Revert Time Count Down      : N/A
Tx Active Change Count      : 1
Last Tx Active Change       : 02/21/2011 13:06:45
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
SAP      : lag-3:1501.1501                        Oper Status: Up
===============================================================================
===============================================================================
```

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.

Label	Description (Continued)
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#
```

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-id* — The ID that displays SAPs for the service in the form *slot*/*mda*/*port*[**.***channel*]. See Common CLI Command Descriptions on page 987 for command syntax.

interface *interface-name* — Displays information for the specified IP interface.

ip-address *ip-address* — Displays information associated with the specified IP address.

detail — Displays detailed information.

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.

Label	Description (Continued)
SAP	The SAP and qtag.
Encap	The encapsulation type of the SAP.
Ethertype	Specifies an Ethernet type II Ethertype value.
Admin State	The administrative state of the SAP.
Oper State	The operating state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapEgressQoSMismatch,RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
Last Status Change	The time of the most recent operating status change to this SAP.
Last Mgmt Change	The time of the most recent management-initiated change to this SAP.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through the port to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the port to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP.
Egress qos-policy	The egress QoS policy ID assigned to the SAP.
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP.
Egress Filter-Id	The egress filter policy ID assigned to the SAP.
Acct. Pol	The accounting policy ID assigned to the SAP.
Collect Stats	Specifies whether collect stats is enabled.
Ignore Oper Down	Displays whether user has enabled or disabled ignore-oper-down parameter.
LLF Admin State	Displays the Link Loss Forwarding administrative state.
LLF Oper State	Displays the Link Loss Forwarding operational state.
Loopback Mode	Displays the Ethernet port loop back mode
Loopback Src Addr	Displays the configured loopback source address

Label	Description (Continued)
Loopback Dst Addr	Displays the configured loopback destination address
No-svc-port used	Displays the port ID of the port on which no service is configured. This port is used for the port loop back with MAC swap functionality.

Sample Output

```
A:Dut-A>config>service>epipe# show service id 2011 sap 1/1/18
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 2011
SAP                : 1/1/18                  Encap             : null
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100
Description        : Default sap description for service id 2011

Admin State        : Up                      Oper State        : Up
Flags              : None
Last Status Change : 07/07/2009 14:39:57
Last Mgmt Change   : 07/07/2009 14:39:14
Admin MTU          : 1514                    Oper MTU          : 1514
LLF Admin State : Up LLF Oper State : Clear
Ingress qos-policy : 10
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
tod-suite          : None
Egr Agg Rate Limit : max                     Endpoint          : N/A

Acct. Pol          : None                    Collect Stats     : Disabled
Ignore Oper Down : Disabled
===============================================================================
A:Dut-A>config>service>epipe#


A:Dut-A>config>service>epipe# show service id 2011 sap 1/1/18 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 2011
SAP                : 1/1/18                  Encap             : null
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100
Description        : Default sap description for service id 2011

Admin State        : Up                      Oper State        : Up
Flags              : None
Last Status Change : 07/07/2009 14:39:57
Last Mgmt Change   : 07/07/2009 14:39:14
Admin MTU          : 1514                    Oper MTU          : 1514
LLF Admin State : Up LLF Oper State : Clear
Ingress qos-policy : 10
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
tod-suite          : None
Egr Agg Rate Limit : max                     Endpoint          : N/A

Acct. Pol          : None                    Collect Stats     : Disabled
```

```
Ignore Oper Down : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress Stats:          0                       0
Egress Stats:           0                       0
-------------------------------------------------------------------------------
Sap per Meter stats
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress Meter 1 (Unicast)
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Meter 2 (Unicast)
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Meter 3 (Unicast)
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Meter 4 (Unicast)
For. InProf           : 0                       0
For. OutProf          : 0                       0
===============================================================================
A:Dut-A>config>service>epipe#

*A:ces-A# show service id 1 sap 1/2/1.1  detail

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/2/1.1                Encap             : cem
Description        : (Not Specified)
Admin State        : Up                     Oper State        : Up
Flags              : None
Last Status Change : 07/06/2010 14:16:41
Last Mgmt Change   : 07/06/2010 11:31:34

Admin MTU          : 1514                   Oper MTU          : 1514
Endpoint           : N/A

Acct. Pol          : None                   Collect Stats     : Disabled
Ignore Oper Down : Disabled
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress Stats:          2815                    613670
Egress Stats:           2815                    613670
-------------------------------------------------------------------------------
CEM SAP Configuration Information
-------------------------------------------------------------------------------
Endpoint Type          : Unstruct. T1     Bit-rate                  : 24
```

```
Payload Size           : 192              Jitter Buffer (ms)         : 5
Jitter Buffer (packets): 6                Playout Threshold (packets): 4
Use RTP Header         : No               Differential               : No
Timestamp Freq         : 0                CAS Framing                : No CAS
Effective PDVT         : +/-2.984 ms

Cfg Alarm      : stray malformed pktloss overrun underrun
Alarm Status   :
-------------------------------------------------------------------------------
CEM SAP Statistics
-------------------------------------------------------------------------------
                          Packets           Seconds           Events
Egress Stats
Forwarded             :  2915
Dropped               :  0
Missing               :  0
Reordered Forwarded   :  0
Underrun              :  0                                   0
Overrun               :  0                                   0
Misordered Dropped    :  0
Malformed Dropped     :  0
LBit Dropped          :  0
Multiple Dropped      :  0
Error                 :                    0
Severely Error        :                    0
Unavailable           :                    0
Failure Count         :                                      0
Jitter Buffer Depth   :  3

Ingress Stats
Forwarded             :  2915
Dropped               :  0
===============================================================================


*A:Dut-A>show# service id 104 sap 1/2/1.2 detail

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 104
SAP                : 1/2/1.2                 Encap             : cem
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Up
Flags              : None
Last Status Change : 12/15/2010 07:39:05
Last Mgmt Change   : 12/15/2010 07:25:37

Admin MTU          : 1514                    Oper MTU          : 1514
Endpoint           : N/A

Acct. Pol          : None                    Collect Stats     : Disabled
Ignore Oper Down : Disabled

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                       Egress qos-policy : 1
-------------------------------------------------------------------------------
Aggregate Policer
-------------------------------------------------------------------------------
```

```
rate               : n/a                       burst             : n/a
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                       Packets                 Octets
Ingress Stats:         786701                  70803090
Egress Stats:          784531                  70607790
Extra-Tag Drop Stats:  n/a                     n/a
-------------------------------------------------------------------------------
CEM SAP Configuration Information
-------------------------------------------------------------------------------
Endpoint Type          : NxDS0           Bit-rate                  : 1
Payload Size           : 64              Jitter Buffer (ms)        : 32
Jitter Buffer (packets): 4               Playout Threshold (packets): 3
Use RTP Header         : No              Differential              : No
Timestamp Freq         : 0               CAS Framing               : No CAS
Effective PDVT         : +/-16.0 ms

Cfg Alarm      : stray malformed pktloss overrun underrun
Alarm Status   :
-------------------------------------------------------------------------------
CEM SAP Statistics
-------------------------------------------------------------------------------
                          Packets          Seconds          Events
Egress Stats
Forwarded             :  784407
Dropped               :  132
Missing               :  0
Reordered Forwarded   :  0
Underrun              :  2355                                1
Overrun               :  0                                   0
Misordered Dropped    :  0
Malformed Dropped     :  0
LBit Dropped          :  132
Multiple Dropped      :  0
Error                 :                  1
Severely Error        :                  0
Unavailable           :                  18
Failure Count         :                                     1
Jitter Buffer Depth   :  2

Ingress Stats
Forwarded             :  786762
Dropped               :  0
===============================================================================
*A:Dut-A>show#

CLI output for 7210 SAS-M and 7210 SAS-T configured in access uplink mode:
*A:SAS-M-A0-2>show>service>id# sap 1/1/1:10.* detail

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/1/1:10.*              Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Up
Flags              : None
Last Status Change : 04/29/2001 06:59:15
Last Mgmt Change   : 04/28/2001 03:09:30
```

```
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100

Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
Admin MTU          : 1522                    Oper MTU          : 1522
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : 1                       Egr Mac Fltr-Id   : n/a
tod-suite          : None
Mac Learning       : Enabled                 Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                 Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled

Acct. Pol          : None                    Collect Stats     : Disabled
Ignore Oper Down : Disabled


-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                      Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Forwarding
Port Number        : 2048                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : N/A
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                     Designated Port   : N/A

Forward transitions: 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                       MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                       Min Auth Interval : 15 minutes

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1
-------------------------------------------------------------------------------
Aggregate Policer
-------------------------------------------------------------------------------
rate               : n/a                     burst             : n/a
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
Classifiers Allocated: 4                   Meters Allocated   : 2
Classifiers Used     : 2                   Meters Used        : 2
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                      Packets                Octets
Ingress Stats:        142761481188           9707780720784
Egress Stats:         0                      0
```

```
Extra-Tag Drop Stats:   n/a                     n/a
-------------------------------------------------------------------------------
Sap per Meter stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Meter 1 (Unicast)
For. InProf           : 17                      1162
For. OutProf          : 0                       0

Ingress Meter 11 (Multipoint)
For. InProf           : 61                      4148
For. OutProf          : 142761547917            9707785259394
===============================================================================
```

sdp

Note : SDP commands are not supported by 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

Syntax **sdp** [*sdp-id* | **far-end** *ip-addr*] [**detail**]

Context show>service>id

Description This command displays information for the SDPs associated with the service.
If no optional parameters are specified, a summary of all associated SDPs is displayed.

Parameters *sdp-id* — 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	The VC type, ether, vlan, or vpls.
VC Tag	The explicit dot1Q value used when encapsulating to the SDP far end.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case).
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The current state of this SDP.

Label	Description (Continued)
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the Keepalive process.
Oper State	The operational state of the Keepalive process.
Hello Time	Transmission frequency of the SDP echo request messages.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	The length of the SDP echo request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field.

Sample Output

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

Admin State       : Up                        Oper State        : Up
Acct. Pol         : None                      Collect Stats     : Disabled
Ingress Label     : 2048                      Egress Label      : 2048
Ing mac Fltr      : n/a                       Egr mac Fltr      : n/a
Ing ip Fltr       : n/a                       Egr ip Fltr       : n/a
```

```
Ing ipv6 Fltr      : n/a                      Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Last Status Change : 05/31/2007 00:45:43      Signaling         : None
Last Mgmt Change   : 05/31/2007 00:45:43
Class Fwding State : Up
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0

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

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

Statistics           :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                        I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                        E. Fwd. Octets    : 0
MCAC Policy Name   :
MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited
Associated LSP LIST :
Lsp Name           : A_B_1
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m35s

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

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

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

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

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

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

Lsp Name           : A_B_8
```

```
Admin State        : Up                        Oper State        : Up
Time Since Last Tr*: 00h26m35s

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

Lsp Name           : A_B_10
Admin State        : Up                        Oper State        : Up
Time Since Last Tr*: 00h26m34s
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : enabled
Default LSP        : A_B_10                    Multicast LSP     : A_B_9
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
af                  A_B_3
be                  A_B_1
ef                  A_B_6
h1                  A_B_7
h2                  A_B_5
l1                  A_B_4
l2                  A_B_2
nc                  A_B_8
===============================================================================
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Up                        Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                       Port State        : Forwarding
Port Number        : 2049                      Port Priority     : 128
Port Path Cost     : 10                        Auto Edge         : Enabled
Admin Edge         : Disabled                  Oper Edge         : N/A
Link Type          : Pt-pt                     BPDU Encap        : Dot1d
Root Guard         : Disabled                  Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                       Designated Port Id: 0
Fwd Transitions    : 0                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                         TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                         RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
* indicates that the corresponding row element may have been truncated.
-------------------------------------------------------------------------------
A:Dut-A#
```

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          : MPLS

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
Class Fwding State : Up
Endpoint          : N/A                       Precedence        : 4
Flags             : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit                 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
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          : MPLS

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
Flags             : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit               Total MAC Addr    : 0
Learned MAC Addr   : 0                      Static MAC Addr   : 0
MAC Learning      : Enabled                 Discard Unkwn Srce: Disabled
MAC Aging         : Enabled
L2PT Termination   : Disabled               BPDU Translation  : Disabled
MAC Pinning       : Disabled
KeepAlive Information :
Admin State       : Disabled                Oper State        : Disabled
Hello Time        : 10                      Hello Msg Len     : 0
Max Drop Count    : 3                       Hold Down Time    : 10
Statistics          :
I. Fwd. Pkts.     : 0                       I. Dro. Pkts.     : 0
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-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          : MPLS
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
Flags             : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
MAC Learning      : Enabled                   Discard Unkwn Srce: Disabled
MAC Aging         : Enabled
L2PT Termination   : Disabled                 BPDU Translation  : Disabled
MAC Pinning       : Disabled
KeepAlive Information :
Admin State       : Disabled                  Oper State        : Disabled
Hello Time        : 10                        Hello Msg Len     : 0
Max Drop Count    : 3                         Hold Down Time    : 10

Statistics            :
I. Fwd. Pkts.     : 0                         I. Dro. Pkts.     : 0
E. Fwd. Pkts.     : 0                         E. Fwd. Octets    : 0
Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
Number of SDPs : 1
===============================================================================
*A:ALA-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          : MPLS
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
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
MAC Learning       : Enabled                  Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
L2PT Termination   : Disabled                 BPDU Translation  : Disabled
MAC Pinning        : Disabled
KeepAlive Information :
Admin State        : Disabled                 Oper State        : Disabled
Hello Time         : 10                       Hello Msg Len     : 0
Max Drop Count     : 3                        Hold Down Time    : 10
Statistics           :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
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#
```

split-horizon-group

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

Context show>service>id

Description This command displays service split horizon groups.

Output

```
*A:7210-SAS>show>service# id 1 split-horizon-group

===============================================================================
Service: Split Horizon Group
===============================================================================
Name                                Description
-------------------------------------------------------------------------------
    access
-------------------------------------------------------------------------------
R = Residential Split Horizon Group
A = Auto Created Split Horizon Group
No. of Split Horizon Groups: 1
===============================================================================
*A:7210-SAS>show>service# id 1 split-horizon-group access

===============================================================================
Service: Split Horizon Group
===============================================================================
Name                                Description
-------------------------------------------------------------------------------
```

```
    access
------------------------------------------------------------------------------
Associations
------------------------------------------------------------------------------
------------------------------------------------------------------------------
R = Residential Split Horizon Group
SAPs Associated : 0                SDPs Associated : 0
*A:7210-SAS>show>service#
```

stp

Syntax **stp** [**detail**]

Context show>service>id

Description This command 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.

Label	Description (Continued)
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.
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.

Label	Description (Continued)
Designated Bridge	Specifies the bridge identifier of the bridge which this port considers to be the designated bridge for this port's segment.

Sample Output

```
A:Dut-A>show>service>id# stp
===============================================================================
Stp info, Service 305
===============================================================================
Bridge Id          : 00:0d.00:20:ab:cd:00:01  Top. Change Count : 5
Root Bridge        : This Bridge              Stp Oper State    : Up
Primary Bridge     : N/A                      Topology Change   : Inactive
Mode               : Rstp                     Last Top. Change  : 0d 08:35:16
Vcp Active Prot.   : N/A
Root Port          : N/A                      External RPC      : 0
===============================================================================
Stp port info
===============================================================================
Sap/Sdp Id         Oper-     Port-      Port-       Port-  Oper-  Link-  Active
                   State     Role       State       Num    Edge   Type   Prot.
-------------------------------------------------------------------------------
1/1/16:305         Up        Designated Forward     2048   False  Pt-pt  Rstp
lag-4:305          Up        Designated Forward     2000   False  Pt-pt  Rstp
1217:305           Up        N/A        Forward     2049   N/A    Pt-pt  N/A
1317:305           Up        N/A        Forward     2050   N/A    Pt-pt  N/A
1417:305           Up        N/A        Forward     2051   N/A    Pt-pt  N/A
1617:305           Pruned    N/A        Discard     2052   N/A    Pt-pt  N/A
===============================================================================
A:Dut-A>show>service>id#


A:Dut-A>show>service>id# stp detail
===============================================================================
Spanning Tree Information
===============================================================================
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up                       Core Connectivity : Down
Stp Admin State    : Up                       Stp Oper State    : Up
Mode               : Rstp                     Vcp Active Prot.  : N/A

Bridge Id          : 00:0d.00:20:ab:cd:00:01  Bridge Instance Id: 13
Bridge Priority    : 0                        Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 0d 08:35:29              Bridge Max Age    : 20
Top. Change Count  : 5                        Bridge Fwd Delay  : 15
MST region revision: 0                        Bridge max hops   : 20
MST region name    :

Root Bridge        : This Bridge
Primary Bridge     : N/A

Root Path Cost     : 0                        Root Forward Delay: 15
Rcvd Hello Time    : 2                        Root Max Age      : 20
Root Priority      : 13                       Root Port         : N/A
-------------------------------------------------------------------------------
```

```
Spanning Tree Sap/Spoke SDP Specifics
--------------------------------------------------------------------------------
SAP Identifier     : 1/1/16:305               Stp Admin State   : Up
Port Role          : Designated               Port State        : Forwarding
Port Number        : 2048                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : PVST
Root Guard         : Disabled                 Active Protocol   : Rstp
Last BPDU from     : 80:04.00:0a:1b:2c:3d:4e
CIST Desig Bridge  : This Bridge              Designated Port   : 34816
Forward transitions: 5                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 29                       RST BPDUs tx      : 23488
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0

SAP Identifier     : lag-4:305                Stp Admin State   : Up
Port Role          : Designated               Port State        : Forwarding
Port Number        : 2000                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Rstp
Last BPDU from     : 80:04.00:0a:1b:2c:3d:4e
CIST Desig Bridge  : This Bridge              Designated Port   : 34768
Forward transitions: 4                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 23                       RST BPDUs tx      : 23454
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0

SDP Identifier     : 1217:305                 Stp Admin State   : Down
Port Role          : N/A                      Port State        : Forwarding
Port Number        : 2049                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last 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 Identifier     : 1317:305                 Stp Admin State   : Down
Port Role          : N/A                      Port State        : Forwarding
Port Number        : 2050                     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 Identifier     : 1417:305                 Stp Admin State   : Down
```

```
Port Role          : N/A                       Port State        : Forwarding
Port Number        : 2051                      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    : 1                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                         TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                         RST BPDUs tx      : 0

SDP Identifier     : 1617:305                  Stp Admin State   : Down
Port Role          : N/A                       Port State        : Discarding
Port Number        : 2052                      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
===============================================================================
A:Dut-A>show>service>id#

*7210-SAS>show>service>id# stp detail

===============================================================================
Spanning Tree Information
===============================================================================

-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up                        Core Connectivity : Down
Stp Admin State    : Up                        Stp Oper State    : Up
Mode               : Mstp                      Vcp Active Prot.  : N/A

Bridge Id          : 80:00.00:25:ba:04:66:a0   Bridge Instance Id: 0
Bridge Priority    : 32768                     Tx Hold Count     : 6
Topology Change    : Inactive                  Bridge Hello Time : 2
Last Top. Change   : 0d 02:54:16               Bridge Max Age    : 20
Top. Change Count  : 27                        Bridge Fwd Delay  : 15

Root Bridge        : 40:00.7c:20:64:ac:ff:63
Primary Bridge     : N/A

Root Path Cost     : 10                        Root Forward Delay: 15
Rcvd Hello Time    : 2                         Root Max Age      : 20
Root Priority      : 16384                     Root Port         : 2048

MSTP info for CIST :
Regional Root      : 80:00.7c:20:64:ad:04:5f   Root Port         : 2048
Internal RPC       : 10                        Remaining Hopcount: 19
MSTP info for MSTI 1 :
Regional Root      : This Bridge               Root Port         : N/A
Internal RPC       : 0                         Remaining Hopcount: 20
```

```
MSTP info for MSTI 2 :
Regional Root      : 00:02.7c:20:64:ad:04:5f  Root Port         : 2048
Internal RPC       : 10                       Remaining Hopcount: 19


-------------------------------------------------------------------------------
Spanning Tree Sap Specifics
-------------------------------------------------------------------------------
SAP Identifier     : 1/1/7:0                  Stp Admin State   : Up
Port Role          : Root                     Port State        : Forwarding
Port Number        : 2048                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Mstp
Last BPDU from     : 80:00.7c:20:64:ad:04:5f  Inside Mst Region : True
CIST Desig Bridge  : 80:00.7c:20:64:ad:04:5f  Designated Port   : 34816
MSTI 1 Port Prio   : 128                      Port Path Cost    : 10
MSTI 1 Desig Brid  : This Bridge              Designated Port   : 34816
MSTI 2 Port Prio   : 128                      Port Path Cost    : 10
MSTI 2 Desig Brid  : 00:02.7c:20:64:ad:04:5f  Designated Port   : 34816
Forward transitions: 17                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 7310                     MST BPDUs tx      : 7277

SAP Identifier     : 1/1/8:0                  Stp Admin State   : Up
Port Role          : Alternate                Port State        : Discarding
Port Number        : 2049                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Mstp
Last BPDU from     : 80:00.7c:20:64:ad:04:5f  Inside Mst Region : True
CIST Desig Bridge  : 80:00.7c:20:64:ad:04:5f  Designated Port   : 34817
MSTI 1 Port Prio   : 128                      Port Path Cost    : 10
MSTI 1 Desig Brid  : This Bridge              Designated Port   : 34817
MSTI 2 Port Prio   : 128                      Port Path Cost    : 10
MSTI 2 Desig Brid  : 00:02.7c:20:64:ad:04:5f  Designated Port   : 34817
Forward transitions: 14                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 7326                     MST BPDUs tx      : 7307

SAP Identifier     : 1/1/9:0                  Stp Admin State   : Up
Port Role          : Designated               Port State        : Forwarding
Port Number        : 2050                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : True
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Mstp
Last BPDU from     : N/A                      Inside Mst Region : True
CIST Desig Bridge  : This Bridge              Designated Port   : 34818
MSTI 1 Port Prio   : 128                      Port Path Cost    : 10
MSTI 1 Desig Brid  : This Bridge              Designated Port   : 34818
MSTI 2 Port Prio   : 128                      Port Path Cost    : 10
MSTI 2 Desig Brid  : This Bridge              Designated Port   : 34818
Forward transitions: 2                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
```

```
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                       MST BPDUs tx      : 7415

SAP Identifier     : 1/1/25:0                Stp Admin State   : Up
Port Role          : Alternate               Port State        : Discarding
Port Number        : 2051                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : False
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : Mstp
Last BPDU from     : 80:00.7c:20:64:ad:04:5f Inside Mst Region : True
CIST Desig Bridge  : 80:00.7c:20:64:ad:04:5f Designated Port   : 34820
MSTI 1 Port Prio   : 128                     Port Path Cost    : 10
MSTI 1 Desig Brid  : This Bridge             Designated Port   : 34819
MSTI 2 Port Prio   : 128                     Port Path Cost    : 10
MSTI 2 Desig Brid  : 00:02.7c:20:64:ad:04:5f Designated Port   : 34820
Forward transitions: 10                      Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 7329                    MST BPDUs tx      : 7303

SAP Identifier     : lag-1:0                 Stp Admin State   : Up
Port Role          : Alternate               Port State        : Discarding
Port Number        : 2052                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : False
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : Mstp
Last BPDU from     : 80:00.7c:20:64:ad:04:5f Inside Mst Region : True
CIST Desig Bridge  : 80:00.7c:20:64:ad:04:5f Designated Port   : 34822
MSTI 1 Port Prio   : 128                     Port Path Cost    : 10
MSTI 1 Desig Brid  : This Bridge             Designated Port   : 34820
MSTI 2 Port Prio   : 128                     Port Path Cost    : 10
MSTI 2 Desig Brid  : 00:02.7c:20:64:ad:04:5f Designated Port   : 34822
Forward transitions: 11                      Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 7322                    MST BPDUs tx      : 7299

===============================================================================
```

VLL 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.

Values service-id: 1 — 214748364
svc-name: A string up to 64 characters in length.

spoke-sdp

Note : SDP commands are not supported by 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

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

ingress-vc-label — Specifies to clear the ingress VC label.

sap

Syntax **sap** *sap-id* {**all** | **cem | counters** | **stp**}

Context clear>service>statistics

Description This command clears SAP statistics for a SAP.

Note: CEM optional parameter is available only on 7210 SAS-M network mode.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 987 for command syntax.

all — Clears all SAP queue statistics and STP statistics.

counters — Clears all queue statistics associated with the SAP.

cem — Clears all CEM statistics associated with the SAP. (CEM optional parameter is available only on 7210 SAS-M network mode)

stp — Clears all STP statistics associated with the SAP.

l2pt — Clears all L2PT statistics associated with the SDP.

cem

Syntax **cem**

Context clear>service>statistics>id

Description Clears the statistics associated with the cpipe service.

sdp

Note : SDP commands are not supported by 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

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 This command clears all traffic queue counters associated with the service ID.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] {**all** | **counters** | **stp**}

Context clear>service>statistics>id

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.

statistics

Syntax **statistics**

Context clear>service

Description This command enables the context to clear statistics for a specific service entity.

VLL 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.

sap

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

Context debug>service>id

Description This command enables debugging for a particular SAP.

Parameters *sap-id* — Specifies the SAP ID.

event-type

Syntax **[no] event-type {arp | config-change |oper-status-change}**

Context debug>service>id

Description This command enables a particular debugging event type.

The **no** form of the command disables the event type debugging.

Parameters **arp —** Displays ARP events.

config-change — Debugs configuration change events.

svc-oper-status-change — Debugs service operational status changes.

Sample Output

```
A:bksim180# debug service id 1000 sap 1/7/1 event-type arp
DEBUG OUTPUT show on CLI is as follows:
3 2008/11/17 18:13:24.35 UTC MINOR: DEBUG #2001 Base Service 1000 SAP 1/7/1 "Service
1000 SAP 1/7/1:
RX: ARP_REQUEST (0x0001)
hwType     : 0x0001
prType     : 0x0800
hwLength   : 0x06
prLength   : 0x04
```

```
srcMac     : 8c:c7:01:07:00:03
destMac    : 00:00:00:00:00:00
srcIp      : 200.1.1.2
destIp     : 200.1.1.1
"

4 2008/11/17 18:13:24.35 UTC MINOR: DEBUG #2001 Base Service 1000 SAP 1/7/1 "Service
1000 SAP 1/7/1:
TX: ARP_RESPONSE (0x0002)
hwType     : 0x0001
prType     : 0x0800
hwLength   : 0x06
prLength   : 0x04
srcMac     : 00:03:0a:0a:0a:0a
destMac    : 8c:c7:01:07:00:03
srcIp      : 200.1.1.1
destIp     : 200.1.1.2
"
```

sdp

Note : SDP commands are not supported by 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>id

Description This command enables debugging for a particular SDP.

Parameters *sdp-id* — Specifies the SDP ID.

VPLS Show Commands

egress-label

Syntax **egress-label** *egress-label1* [*egress-label2*]

Context show>service

Description This command displays service information using the range of egress labels.

If only the mandatory *egress-label1* parameter is specified, only services using the specified label are displayed.

If both *egress-label1* and *egress-label2* parameters are specified, the services using the range of labels X where *egress-label1* <= X <= *egress-label2* are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *egress-label1* — 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

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 value that identifies a service.
Mac Move	Indicates the administrative state of the MAC movement feature associated with the service.

Label	Description (Continued)
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. 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.
Total Count	The current number of entries (both learned and static) in the FDB of this service.
Learned Count	The current number of learned entries in the FDB of this service.
Static Count	The current number of static entries in the FDB of this service.
Remote Age	The number of seconds used to age out FDB entries learned on an SDP. These entries correspond to MAC addresses learned on remote SAPs.
Local Age	The seconds used to age out FDB entries learned on local SAPs.
High WaterMark	The utilization of the FDB table of this service at which a 'table full' alarm is raised by the agent.
Low WaterMark	The utilization of the FDB table of this service at which a 'table full' alarm is cleared by the agent.
Mac Learning	Specifies whether the MAC learning process is enabled in this service.
Discard Unknown	Specifies whether frames received with an unknown destination MAC are discarded in this service.
MAC Aging	Specifies whether the MAC aging process is enabled in this service.
MAC Pinning	Specifies whether MAC pinning is enabled in this service.
Relearn Only	When enabled, indicates that either the FDB table of this service is full or that the maximum system-wide number of 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 FDB	The current number of service FDBs configured on this node.
Total FDB Configured Size	The sum of configured FDBs.

Label	Description (Continued)
Total FDB Entries In Use	The total number of entries (both learned and static) in use.

Sample Output

```
A:7210-SASE# show service fdb-info
===============================================================================
Forwarding Database(FDB) Information
===============================================================================
Service Id        : 1                   Mac Move          : Disabled
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Table Size        : 8191                Total Count       : 675
Learned Count     : 675                 Static Count      : 0
Local Age         : 60
High WaterMark    : 5%                  Low Watermark     : 1%
Mac Learning      : Enabl               Discard Unknown   : Dsabl
Mac Aging         : Enabl               Relearn Only      : False

Service Id        : 2                   Mac Move          : Disabled
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Table Size        : 8191                Total Count       : 0
Learned Count     : 0                   Static Count      : 0
Local Age         : 80
High WaterMark    : 10%                 Low Watermark     : 2%
Mac Learning      : Enabl               Discard Unknown   : Dsabl
Mac Aging         : Enabl               Relearn Only      : False

Service Id        : 3                   Mac Move          : Disabled
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Table Size        : 8191                Total Count       : 675
Learned Count     : 675                 Static Count      : 0
Local Age         : 100
High WaterMark    : 15%                 Low Watermark     : 3%
Mac Learning      : Enabl               Discard Unknown   : Dsabl
Mac Aging         : Enabl               Relearn Only      : False

Service Id        : 4                   Mac Move          : Disabled
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Table Size        : 8191                Total Count       : 0
Learned Count     : 0                   Static Count      : 0
Local Age         : 120
High WaterMark    : 20%                 Low Watermark     : 4%
Mac Learning      : Enabl               Discard Unknown   : Dsabl
Mac Aging         : Enabl               Relearn Only      : False

Service Id        : 5                   Mac Move          : Disabled
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Table Size        : 8191                Total Count       : 0
Learned Count     : 0                   Static Count      : 0
Local Age         : 600
High WaterMark    : 25%                 Low Watermark     : 5%
Mac Learning      : Enabl               Discard Unknown   : Dsabl
Mac Aging         : Enabl               Relearn Only      : False

Service Id        : 6                   Mac Move          : Disabled
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Table Size        : 8191                Total Count       : 675
```

```
Learned Count     : 675                Static Count      : 0
Local Age         : 86400
High WaterMark    : 30%                Low Watermark     : 10%
Mac Learning      : Enabl              Discard Unknown   : Dsabl
Mac Aging         : Enabl              Relearn Only      : False
-------------------------------------------------------------------------------
Total Service FDBs : 6
Total FDB Configured Size : 49146
Total FDB Entries In Use  : 2025
-------------------------------------------------------------------------------
===============================================================================
A:7210-SASE#
```

fdb-mac

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

Context show>service

Description This command 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 where the MAC is defined.
Type/Age	Static – FDB entries created by management. Learned – Dynamic entries created by the learning process. OAM – Entries created by the OAM process. H – Host, the entry added by the system for a static configured subscriber host. D or DHCP – DHCP-installed MAC. Learned addresses can be temporarily frozen by the DHCP snooping application for the duration of a DHCP lease. P — Indicates the MAC is protected by the MAC protection feature.

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 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 spoke.
I.Lbl	The ingress label used by the far-end device to send packets to this device in this service by the SDP.
E.Lbl	The egress label used by this device to send packets to the far-end device in this service by the SDP.
Number of Bindings Found	The number of SDP bindings within the label range specified.

sap-using

Syntax **sap-using interface** [*ip-address* | *ip-int-name*]
sap-using [**ingress** | **egress**] **filter** *filter-id*
sap-using [**sap** *sap-id*]
sap-using [**ingress**] **qos-policy** *qos-policy-id*

Context show>service

Description This command displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

Parameters **ingress —** Specifies matching an ingress policy.

egress — Specifies matching an egress policy.

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

Values 1 — 65535

sap-id — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 987 for command syntax.

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.
I.QoS	The SAP ingress QoS policy number specified on the ingress SAP.
I.MAC/IP	The MAC or IP filter policy ID applied to the ingress SAP.
Egr. Fltr	The filter policy ID applied to the egress SAP.
A.Pol	The accounting policy ID assigned to the SAP.
Adm	The administrative state of the SAP.
Opr	The actual state of the SAP.

Sample Output

```
*A:ALU_SIM2>config>service>vpls# show service sap-using
===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.   Adm  Opr
                                           QoS   Fltr    Fltr
-------------------------------------------------------------------------------
1/1/1:10                        1          1     none     none   Up   Up
1/1/3:500.*                     1          1     none     none   Up   Up
```

```
1/1/1:200                       200        1     none     none   Up   Up
1/1/3:100.200                   200        1     none     none   Up   Up
1/1/1:300                       300        1     none     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 5
-------------------------------------------------------------------------------
*A:ALU_SIM2>config>service>vpls#
```

sdp

Note : SDP commands are not supported by 7210 SAS-M devices configured in uplink mode.

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

Context show>service>id

Description This command displays information for the SDPs associated with the service.

If no optional parameters are specified, a summary of all associated SDPs is displayed.

Parameters *sdp-id* — Displays only information for the specified SDP ID. An SDP is a logical mechanism that ties a far-end 7210 SAS M to a particular service without having to specifically define far end SAPs. Each SDP represents a method to reach a 7210 SAS M router.

Default All SDPs.

Values 1 — 17407

far-end *ip-addr* **—** Displays only SDPs matching with the specified system IP address of the far-end destination 7210 SAS M router for the Service Distribution Point (SDP) that is the termination point for a service.

Default SDPs with any far-end IP address.

detail — Displays detailed SDP information.

Output **Show Service 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.
VC Type	Displays the VC type, ether or vlan.
VC Tag	Displays the explicit dot1Q value used when encapsulating to the SDP far end.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case.)

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 MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the Keepalive process.
Oper State	The operational state of the Keepalive process.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
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.

sdp-using

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

Context show>service

Description This command displays services using SDP or far-end address options.

Parameters *sdp-id* — Displays only services bound to the specified SDP ID.

Values 1 — 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 —** The following table describes service-using output fields.

Label	Description
Svc ID	The service identifier.
Sdp ID	The SDP identifier.
Type	Specifies the type of SDP: Spoke.
Far End	The far-end address of the SDP.
Oper State	The operational state of the service.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.

Sample Output

```
*A:ALA-1# show service sdp-using 300
===============================================================================
Service Destination Point (Sdp Id : 300)
===============================================================================
SvcId       SdpId               Type Far End        Opr State I.Label  E.Label
-------------------------------------------------------------------------------
2           300:2               Spok 10.0.0.13      Up        131070   131070
-------------------------------------------------------------------------------
Number of SDPs : 51
-------------------------------------------------------------------------------
*A:ALA-1#
```

service-using

Syntax **service-using** [**epipe**] [**vpls**] [**mirror**] [**customer** *customer-id*]

Context show>service

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

Parameters **epipe —** Displays matching Epipe services.

vpls — Displays matching VPLS instances.

mirror — Displays matching mirror services.

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

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 administrative state of the service.
Opr	The operating state of the service.
CustomerID	The ID of the customer who owns this service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this service.

Sample Output

```
*A:ALA-12# show service service-using customer 10
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
1            VPLS      Up     Up         10                09/05/2006 13:24:15
100          IES       Up     Up         10                09/05/2006 13:24:15
300          Epipe     Up     Up         10                09/05/2006 13:24:15
-------------------------------------------------------------------------------
Matching Services : 3
===============================================================================
*A:ALA-12#


*A:ALA-12# show service service-using epipe
===============================================================================
```

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


*A:ALA-14# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
10           mVPLS     Down   Down       1                 10/26/2006 15:44:57
11           mVPLS     Down   Down       1                 10/26/2006 15:44:57
100          mVPLS     Up     Up         1                 10/26/2006 15:44:57
101          mVPLS     Up     Up         1                 10/26/2006 15:44:57
102          mVPLS     Up     Up         1                 10/26/2006 15:44:57
-------------------------------------------------------------------------------
Matching Services : 5
-------------------------------------------------------------------------------
*A:ALA-14#



A:Dut-A>config>service# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
100          mVPLS     Up     Up         1                 07/07/2009 14:39:13
101          uVPLS     Up     Up         1                 07/07/2009 14:39:13
102          uVPLS     Up     Up         1                 07/07/2009 14:39:13
103          uVPLS     Up     Up         1                 07/07/2009 14:39:13
104          uVPLS     Up     Up         1                 07/07/2009 14:39:13
105          uVPLS     Up     Up         1                 07/07/2009 14:39:13
201          VPLS      Up     Up         1                 07/07/2009 14:39:13
202          VPLS      Up     Up         1                 07/07/2009 14:39:13
203          VPLS      Up     Up         1                 07/07/2009 14:39:13
204          VPLS      Up     Up         1                 07/07/2009 14:39:13
205          VPLS      Up     Up         1                 07/07/2009 14:39:13
300          mVPLS     Up     Up         1                 07/07/2009 14:39:13
301          uVPLS     Up     Up         1                 07/07/2009 14:39:13
302          uVPLS     Up     Up         1                 07/07/2009 14:39:13
303          uVPLS     Up     Up         1                 07/07/2009 14:39:13
304          uVPLS     Up     Up         1                 07/07/2009 14:39:1
305          uVPLS     Up     Up         1                 07/07/2009 14:39:1
401          VPLS      Up     Up         1                 07/07/2009 14:39:1
402          VPLS      Up     Up         1                 07/07/2009 14:39:1
403          VPLS      Up     Up         1                 07/07/2009 14:39:1
404          VPLS      Up     Up         1                 07/07/2009 14:39:1
405          VPLS      Up     Up         1                 07/07/2009 14:39:1
```

```
500          mVPLS     Up     Up         1                  07/07/2009 14:39:1
511          uVPLS     Up     Up         1                  07/07/2009 14:39:1
513          uVPLS     Up     Up         1                  07/07/2009 14:39:1
515          uVPLS     Up     Up         1                  07/07/2009 14:39:1
517          uVPLS     Up     Up         1                  07/07/2009 14:39:1
519          uVPLS     Up     Up         1                  07/07/2009 14:39:1
601          VPLS      Up     Up         1                  07/07/2009 14:39:1
602          VPLS      Up     Up         1                  07/07/2009 14:39:1
603          VPLS      Up     Up         1                  07/07/2009 14:39:1
604          VPLS      Up     Up         1                  07/07/2009 14:39:1
605          VPLS      Up     Up         1                  07/07/2009 14:39:1
701          VPLS      Up     Up         1                  07/07/2009 14:39:1
702          VPLS      Up     Up         1                  07/07/2009 14:39:1
703          VPLS      Up     Up         1                  07/07/2009 14:39:1
704          VPLS      Up     Up         1                  07/07/2009 14:39:1
801          VPLS      Up     Up         1                  07/07/2009 14:39:1
802          VPLS      Up     Up         1                  07/07/2009 14:39:1
803          VPLS      Up     Up         1                  07/07/2009 14:39:1
804          VPLS      Up     Up         1                  07/07/2009 14:39:1
805          VPLS      Up     Up         1                  07/07/2009 14:39:1
901          VPLS      Up     Up         1                  07/07/2009 14:39:1
902          VPLS      Up     Up         1                  07/07/2009 14:39:1
903          VPLS      Up     Up         1                  07/07/2009 14:39:1
904          VPLS      Up     Up         1                  07/07/2009 14:39:1
905          VPLS      Up     Up         1                  07/07/2009 14:39:1
906          VPLS      Up     Up         1                  07/07/2009 14:39:1
907          VPLS      Up     Up         1                  07/07/2009 14:39:1
908          VPLS      Up     Up         1                  07/07/2009 14:39:1
909          VPLS      Up     Up         1                  07/07/2009 14:39:1
910          VPLS      Up     Up         1                  07/07/2009 14:39:1
1101         Epipe     Up     Up         1                  07/07/2009 14:39:1
1102         Epipe     Up     Up         1                  07/07/2009 14:39:1
1103         Epipe     Up     Up         1                  07/07/2009 14:39:1
1104         Epipe     Up     Up         1                  07/07/2009 14:39:1
1105         Epipe     Up     Up         1                  07/07/2009 14:39:1
1501         Epipe     Up     Up         1                  07/07/2009 14:39:1
1502         Epipe     Up     Up         1                  07/07/2009 14:39:1
1503         Epipe     Up     Up         1                  07/07/2009 14:39:1
1504         Epipe     Up     Up         1                  07/07/2009 14:39:1
1505         Epipe     Up     Up         1                  07/07/2009 14:39:1
2001         Mirror    Up     Up         1                  07/07/2009 14:39:1
2002         Mirror    Up     Up         1                  07/07/2009 14:39:1
2011         Epipe     Up     Up         1                  07/07/2009 14:39:1
2012         VPLS      Up     Up         1                  07/07/2009 14:39:1
3000         mVPLS     Up     Up         1                  07/07/2009 14:39:1
4001         VPLS      Up     Up         1                  07/07/2009 14:39:1
4002         VPLS      Up     Up         1                  07/07/2009 14:39:1
-------------------------------------------------------------------------------
Matching Services : 69
===============================================================================
A:Dut-A>config>service#
```

id

Syntax **id** *service-id*

Context show>service

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

Parameters *service-id* — The unique service identification number that identifies the service in the service domain.

Values service-id: 1 — 214748364
svc-name: A string up to 64 characters in length.

all — Display detailed information about the service.

base — Display basic service information.

endpoint — Display service endpoint information.

fdb — Display FDB entries.

labels — Display labels being used by this service.

mstp-configuration — - Display MSTP information.

sap — Display SAPs associated to the service.

sdp — Display SDPs associated with the service.

stp — Display STP information.

all

Syntax **all**

Context show>service>id

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

Output **Show service ID all output —** The following table describes the command output fields.

Label	Description
Service Id	The service identifier.
VPN Id	The number which identifies the VPN.
Service Type	Specifies the type of service.
SDP Id	The SDP identifier.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.

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	Name of the split horizon group for this service.
Description	Description of the split horizon group.
Last Changed	The date and time of the most recent management-initiated change to this split horizon group.
SDP Id	The SDP identifier.
Type	Indicates whether this service SDP binding is a spoke or a mesh.
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
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.

Label	Description (Continued)
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field.
Number of SDPs	The total number SDPs applied to this service ID.
Service Id	The service identifier.
Port Id	The ID of the access port where this SAP is defined.
Description	Generic information about the SAP.
Encap Value	The value of the label used to identify this SAP on the access port.
Admin State	The administrative state of the SAP.
Oper State	The operating state of the SAP.
Last Changed	The date and time of the last change.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The SAP ingress QoS policy ID.
Ingress Filter-Id	The SAP ingress filter policy ID.
Egress Filter-Id	The SAP egress filter policy ID.
Multi Svc Site	Indicates the multi-service site that the SAP is a member.
Ingress sched-policy	Indicates the ingress QoS scheduler for the SAP.
Egress sched-policy	Indicates the egress QoS scheduler for the SAP.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
Ingress Stats	The number of received packets/octets for this SAP.
Egress Stats	The number of packets/octets forwarded out of this SAP.
Ingress Meter 1	The index of the ingress QoS meter of this SAP.
High priority offered	The packets or octets count of the high priority traffic for the SAP.

Label	Description (Continued)
For.InProf	The packets or octets count of the in-profile forwarded traffic for the SAP.
For.OutProf	The number of out of profile traffic packets/octets forwarded.
Managed by Service	Specifies the service-id of the management VPLS managing this SAP.
Managed by MSTI	Specifies the MST instance inside the management VPLS managing this SAP.
Last BPDU from	The bridge ID of the sender of the last BPDU received on this SAP.
Managed by SAP	Specifies the sap-id inside the management VPLS managing this SAP.
Prune state	Specifies the STP state inherited from the management VPLS.
Managed by Service	Specifies the service-id of the management VPLS managing this 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:Dut-A>config>service#  show service id 305 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 305                 Vpn Id            : 305
Service Type      : uVPLS
Description       : Default tls description for service id 305
Customer Id       : 1
Last Status Change: 07/07/2009 14:39:57
Last Mgmt Change  : 07/07/2009 14:39:14
Admin State       : Up                  Oper State        : Up
MTU               : 1514
MTU Check         : Disabled
SAP Count         : 2                   SDP Bind Count    : 4
Send Flush on Fail: Disabled
Uplink Type       : MPLS
Propagate MacFlush: Disabled
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
 Sdp Id 1217:305  -(10.20.1.2)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 1217:305                 Type              : Spoke
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 9186
Far End            : 10.20.1.2                Delivery          : MPLS

Admin State       : Up                       Oper State        : Up
```

```
Acct. Pol          : None                     Collect Stats     : Disabled
Managed by Service : 300                      Prune State       : Not Pruned
Managed by Spoke   : 1217:300
Ingress Label      : 130506                   Egress Label      : 130516
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Last Status Change : 07/07/2009 18:49:40      Signaling         : TLDP
Last Mgmt Change   : 07/07/2009 14:39:14      Force Vlan-Vc     : Disabled
Last Mgmt Change   : 07/07/2009 14:39:14
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0

MAC Learning       : Enabled                  Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
L2PT Termination   : Disabled                 BPDU Translation  : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False                    Block On Mesh Fail: False

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

Statistics           :
I. Fwd. Pkts.      : 13601                    I. Fwd. Octs.     : 10676338
E. Fwd. Pkts.      : 65165676                 E. Fwd. Octets    : 39462444830

Associated LSP LIST :
Lsp Name           : A_B_17
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 05h24m26s
-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Down                     Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Forwarding
Port Number        : 2049                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                      Designated Port Id: 0

Fwd Transitions    : 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
-------------------------------------------------------------------------------
 Sdp Id 1317:305  -(10.20.1.3)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 1317:305                 Type              : Spoke
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 9186
Far End            : 10.20.1.3                Delivery          : MPLS
```

```
Admin State        : Up                       Oper State        : Up
Acct. Pol          : None                     Collect Stats     : Disabled
Managed by Service : 300                      Prune State       : Not Pruned
Managed by Spoke   : 1317:300
Ingress Label      : 130454                   Egress Label      : 130591
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Last Status Change : 07/07/2009 18:49:43      Signaling         : TLDP
Last Mgmt Change   : 07/07/2009 14:39:14      Force Vlan-Vc     : Disabled
Last Mgmt Change   : 07/07/2009 14:39:14
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0

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

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

Statistics           :
I. Fwd. Pkts.      : 10100                    I. Fwd. Octs.     : 7178960
E. Fwd. Pkts.      : 65466629                 E. Fwd. Octets    : 39665246044

Associated LSP LIST :
Lsp Name           : A_C_17
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 05h24m23s
-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Down                     Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Forwarding
Port Number        : 2050                     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 1417:305  -(10.20.1.4)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 1417:305                 Type              : Spoke
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 9186
Far End            : 10.20.1.4                Delivery          : MPLS
```

```
Admin State        : Up                       Oper State        : Up
Acct. Pol          : None                     Collect Stats     : Disabled
Managed by Service : 300                      Prune State       : Not Pruned
Managed by Spoke   : 1417:300
Ingress Label      : 130428                   Egress Label      : 131015
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Last Status Change : 07/07/2009 18:13:42      Signaling         : TLDP
Last Mgmt Change   : 07/07/2009 14:39:14      Force Vlan-Vc     : Disabled
Last Mgmt Change   : 07/07/2009 14:39:14
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 250
Learned MAC Addr   : 250                      Static MAC Addr   : 0

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

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

Statistics            :
I. Fwd. Pkts.      : 97516328                 I. Fwd. Octs.     : 47531982212
E. Fwd. Pkts.      : 166191635                E. Fwd. Octets    : 67215031404

Associated LSP LIST :
Lsp Name           : A_D_17
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 09h33m18s
-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Down                     Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Forwarding
Port Number        : 2051                     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    : 1                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
-------------------------------------------------------------------------------
 Sdp Id 1617:305  -(10.20.1.6)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 1617:305                 Type              : Spoke
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 9186
Far End            : 10.20.1.6                Delivery          : MPLS
```

```
Admin State        : Up                       Oper State        : Up
Acct. Pol          : None                     Collect Stats     : Disabled
Managed by Service : 300                      Prune State       : Pruned
Managed by Spoke   : 1617:300
Ingress Label      : 131060                   Egress Label      : 130843
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Last Status Change : 07/07/2009 14:40:52      Signaling         : TLDP
Last Mgmt Change   : 07/07/2009 14:39:14      Force Vlan-Vc     : Disabled
Last Mgmt Change   : 07/07/2009 14:39:14
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0

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

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

Statistics           :
I. Fwd. Pkts.      : 12889                    I. Fwd. Octs.     : 6000654
E. Fwd. Pkts.      : 11999                    E. Fwd. Octets    : 5208494

Associated LSP LIST :
Lsp Name           : A_F_17
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 09h33m18s
-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Down                     Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Discarding
Port Number        : 2052                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                      Designated Port Id: 0

Fwd Transitions    : 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 4
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
SAP 1/1/16:305
-------------------------------------------------------------------------------
Service Id         : 305
SAP                : 1/1/16:305               Encap             : q-tag
```

```
Dot1Q Ethertype    : 0x8100                    QinQ Ethertype    : 0x8100
Description        : Default sap description for service id 305

Admin State        : Up                        Oper State        : Up
Flags              : None
Last Status Change : 07/07/2009 14:39:57
Last Mgmt Change   : 07/07/2009 14:39:14
Max Nbr of MAC Addr: No Limit                  Total MAC Addr    : 0
Learned MAC Addr   : 0                         Static MAC Addr   : 0
Admin MTU          : 9212                      Oper MTU          : 9212
Ingress qos-policy : 10
Ingr IP Fltr-Id    : n/a                       Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : 305                       Egr Mac Fltr-Id   : n/a
tod-suite          : None
Egr Agg Rate Limit : max
Mac Learning       : Enabled                   Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                   Mac Pinning       : Disabled
L2PT Termination   : Disabled                  BPDU Translation  : Disabled

Acct. Pol          : None                      Collect Stats     : Disabled
-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Up                        Stp Oper State    : Up
Core Connectivity  : Down
Port Role          : Designated                Port State        : Forwarding
Port Number        : 2048                      Port Priority     : 128
Port Path Cost     : 10                        Auto Edge         : Enabled
Admin Edge         : Disabled                  Oper Edge         : False
Link Type          : Pt-pt                     BPDU Encap        : Dot1d
Root Guard         : Disabled                  Active Protocol   : Rstp
Last BPDU from     : 80:04.00:0a:1b:2c:3d:4e
CIST Desig Bridge  : This Bridge               Designated Port   : 34816

Forward transitions: 5                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                         TCN BPDUs tx      : 0
RST BPDUs rcvd     : 29                        RST BPDUs tx      : 17610
MST BPDUs rcvd     : 0                         MST BPDUs tx      : 0
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                         Packets                 Octets
Ingress Stats:           66655                   39685976
Egress Stats:            65864342                38651746348
-------------------------------------------------------------------------------
Sap per Meter stats
-------------------------------------------------------------------------------
                         Packets                 Octets

Ingress Meter 1 (Unicast)
For. InProf           : 0                        0
For. OutProf          : 0                        0

Ingress Meter 2 (Unicast)
For. InProf           : 0                        0
For. OutProf          : 0                        0

Ingress Meter 3 (Unicast)
For. InProf           : 0                        0
```

```
For. OutProf          : 0                       0

Ingress Meter 4 (Unicast)
For. InProf           : 11406                   4291328
For. OutProf          : 12575                   4325376

Ingress Meter 11 (Multipoint)
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Meter 12 (Multipoint)
For. InProf           : 3108                    3108000
For. OutProf          : 2235                    2235000

Ingress Meter 13 (Multipoint)
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Meter 14 (Multipoint)
For. InProf           : 8772                    5166272
For. OutProf          : 4840                    3072000
-------------------------------------------------------------------------------
SAP lag-4:305
-------------------------------------------------------------------------------
Service Id         : 305
SAP                : lag-4:305                Encap             : q-tag
Description        : Default sap description for service id 305

Admin State        : Up                       Oper State        : Up
Flags              : None
Last Status Change : 07/07/2009 14:39:57
Last Mgmt Change   : 07/07/2009 14:39:14
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 125
Learned MAC Addr   : 125                      Static MAC Addr   : 0
Admin MTU          : 9212                     Oper MTU          : 9212
Ingress qos-policy : 10
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : 305                      Egr Mac Fltr-Id   : n/a
tod-suite          : None
Egr Agg Rate Limit : max
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
L2PT Termination   : Disabled                 BPDU Translation  : Disabled

Acct. Pol          : None                     Collect Stats     : Disabled
-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Up                       Stp Oper State    : Up
Core Connectivity  : Down
Port Role          : Designated               Port State        : Forwarding
Port Number        : 2000                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Rstp
Last BPDU from     : 80:04.00:0a:1b:2c:3d:4e
CIST Desig Bridge  : This Bridge              Designated Port   : 34768

Forward transitions: 4                        Bad BPDUs rcvd    : 0
```

```
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 23                       RST BPDUs tx      : 17578
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                       Packets                  Octets
Ingress Stats:         190824363                87464904956
Egress Stats:          97572636                 45409567760
-------------------------------------------------------------------------------
Sap per Meter stats
-------------------------------------------------------------------------------
                       Packets                  Octets
Ingress Meter 1 (Unicast)
For. InProf          : 0                        0
For. OutProf         : 0                        0

Ingress Meter 2 (Unicast)
For. InProf          : 0                        0
For. OutProf         : 0                        0

Ingress Meter 3 (Unicast)
For. InProf          : 0                        0
For. OutProf         : 0                        0

Ingress Meter 4 (Unicast)
For. InProf          : 56963244                 20851041536
For. OutProf         : 59512115                 19403302144

Ingress Meter 11 (Multipoint)
For. InProf          : 0                        0
For. OutProf         : 0                        0

Ingress Meter 12 (Multipoint)
For. InProf          : 12922550                 12922550000
For. OutProf         : 9452800                  9452800000

Ingress Meter 13 (Multipoint)
For. InProf          : 0                        0
For. OutProf         : 0                        0

Ingress Meter 14 (Multipoint)
For. InProf          : 43268112                 21539479708
For. OutProf         : 6788456                  2546422464
-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up                       Core Connectivity : Down
Stp Admin State    : Up                       Stp Oper State    : Up
Mode               : Rstp                     Vcp Active Prot.  : N/A

Bridge Id          : 00:0d.00:20:ab:cd:00:01  Bridge Instance Id: 13
Bridge Priority    : 0                        Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 0d 05:21:37              Bridge Max Age    : 20
Top. Change Count  : 5                        Bridge Fwd Delay  : 15
MST region revision: 0                        Bridge max hops   : 20
MST region name    :

Root Bridge        : This Bridge
```

```
Primary Bridge     : N/A

Root Path Cost     : 0                        Root Forward Delay: 15
Rcvd Hello Time    : 2                        Root Max Age      : 20
Root Priority      : 13                       Root Port         : N/A
-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id         : 305                Mac Move          : Disabled
Mac Move Rate      : 2                  Mac Move Timeout  : 10
Table Size         : 500                Total Count       : 375
Learned Count      : 375                Static Count      : 0
Remote Age         : 60                 Local Age         : 60
High WaterMark     : 95%                Low Watermark     : 90%
Mac Learning       : Enabl              Discard Unknown   : Dsabl
Mac Aging          : Enabl              Relearn Only      : False
===============================================================================
A:Dut-A>config>service#
```

Sample output for 7210 SAS-M in access uplink mode:

```
*A:SAS-M-A0-2>show>service>id# all

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

Service Id         : 1                  Vpn Id            : 0
Service Type       : VPLS
Description        : (Not Specified)
Customer Id        : 1
Last Status Change: 04/29/2001 06:59:15
Last Mgmt Change   : 04/28/2001 03:03:03
Admin State        : Up                 Oper State        : Up
MTU                : 1514
MTU Check          : Enabled
SAP Count          : 2                  SDP Bind Count    : 0
Snd Flush on Fail  : Disabled
Uplink Type:       : MPLS
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/1/1:10.*
-------------------------------------------------------------------------------
Service Id         : 1
SAP                : 1/1/1:10.*             Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                     Oper State        : Up
Flags              : None
Last Status Change : 04/29/2001 06:59:15
Last Mgmt Change   : 04/28/2001 03:09:30
Dot1Q Ethertype    : 0x8100                 QinQ Ethertype    : 0x8100

Max Nbr of MAC Addr: No Limit               Total MAC Addr    : 0
Learned MAC Addr   : 0                      Static MAC Addr   : 0
```

```
Admin MTU          : 1522                    Oper MTU          : 1522
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : 1                       Egr Mac Fltr-Id   : n/a
tod-suite          : None
Mac Learning       : Enabled                 Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                 Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled

Acct. Pol          : None                    Collect Stats     : Disabled



-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                      Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Forwarding
Port Number        : 2048                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : N/A
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                     Designated Port   : N/A

Forward transitions: 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                       MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                       Min Auth Interval : 15 minutes

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1
-------------------------------------------------------------------------------
Aggregate Policer
-------------------------------------------------------------------------------
rate               : n/a                     burst             : n/a
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
Classifiers Allocated: 4                    Meters Allocated   : 2
Classifiers Used     : 2                    Meters Used        : 2
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                       Packets                 Octets
Ingress Stats:         142761481188            9707780720784
Egress Stats:          0                       0
Extra-Tag Drop Stats:  n/a                     n/a
-------------------------------------------------------------------------------
Sap per Meter stats
-------------------------------------------------------------------------------
```

```
                      Packets                 Octets

Ingress Meter 1 (Unicast)
For. InProf           : 17                    1162
For. OutProf          : 0                     0

Ingress Meter 11 (Multipoint)
For. InProf           : 61                    4148
For. OutProf          : 142761547917          9707785259394

-------------------------------------------------------------------------------
SAP 1/1/2:10.*
-------------------------------------------------------------------------------
Service Id         : 1
SAP                : 1/1/2:10.*              Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Up
Flags              : None
Last Status Change : 04/29/2001 07:03:49
Last Mgmt Change   : 04/28/2001 03:02:15
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100

Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
Admin MTU          : 1522                    Oper MTU          : 1522
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
tod-suite          : None
Mac Learning       : Enabled                 Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                 Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled

Acct. Pol          : None                    Collect Stats     : Disabled



-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                      Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Forwarding
Port Number        : 2049                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : N/A
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                     Designated Port   : N/A

Forward transitions: 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                       MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
```

```
Host Limit         : 1                        Min Auth Interval : 15 minutes

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1
-------------------------------------------------------------------------------
Aggregate Policer
-------------------------------------------------------------------------------
rate               : n/a                      burst             : n/a
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
Classifiers Allocated: 4                    Meters Allocated   : 2
Classifiers Used     : 2                    Meters Used        : 2
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                       Packets                 Octets
Ingress Stats:         0                       0
Egress Stats:          535194841               36393249188
Extra-Tag Drop Stats:  n/a                     n/a
-------------------------------------------------------------------------------
Sap per Meter stats
-------------------------------------------------------------------------------
                       Packets                 Octets

Ingress Meter 1 (Unicast)
For. InProf          : 0                       0
For. OutProf         : 0                       0

Ingress Meter 11 (Multipoint)
For. InProf          : 0                       0
For. OutProf         : 0                       0

-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up                      Core Connectivity : Down
Stp Admin State    : Down                    Stp Oper State    : Down
Mode               : Rstp                    Vcp Active Prot.  : N/A

Bridge Id          : 80:00.00:25:ba:02:ea:00 Bridge Instance Id: 0
Bridge Priority    : 32768                   Tx Hold Count     : 6
Topology Change    : Inactive                Bridge Hello Time : 2
Last Top. Change   : 0d 00:00:00             Bridge Max Age    : 20
Top. Change Count  : 0                       Bridge Fwd Delay  : 15

Root Bridge        : N/A
Primary Bridge     : N/A

Root Path Cost     : 0                       Root Forward Delay: 15
Rcvd Hello Time    : 2                       Root Max Age      : 20
Root Priority      : 32768                   Root Port         : N/A

-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id        : 1                  Mac Move          : Disabled
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Mac Move Retries  : 3
```

```
Table Size        : 250                Total Count       : 0
Learned Count     : 0                  Static Count      : 0
Remote Age        : 900                Local Age         : 300
High Watermark    : 95%                Low Watermark     : 90%
Mac Learning      : Enabled            Discard Unknown   : Disabled
Mac Aging         : Enabled            Relearn Only      : False

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:SAS-M-A0-2>show>service>id#
```

arp

Syntax **arp** [*ip-address*] | [**mac** *ieee-address*] | [**sap** *sap-id*] | [**interface** *ip-int-name*]

Context show>service>id

Description This command displays the ARP table for the VPLS instance. The ARP entries for a subscriber interface are displayed uniquely. Each MAC associated with the subscriber interface child group-interfaces is displayed with each subscriber interface ARP entry for easy lookup.

Parameters *ip-address* — All IP addresses.

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

Default All MAC addresses.

sap *sap-id* **—** Displays SAP information for the specified SAP ID.

interface — Specifies matching service ARP entries associated with the IP interface.

ip-address — 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
IP Address	The IP address.
MAC Address	The specified MAC address.
	Type Static — FDB entries created by management.
	Learned — Dynamic entries created by the learningprocess.
	Other — Local entries for the IP interfaces created.

Label	Description
Expiry	The age of the ARP entry.
Interface	The interface applied to the service.
SAP	The SAP ID.

base

Syntax **base** [**msap**]

Context show>service>id
show>service>id>igmp-snooping

Description This command displays basic information about the service ID including service type, description, SAPs and SDP.

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

Label	Description
Service Id	The service identifier.
Service Type	Displays the type of service.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Adm	The administrative state of the service.
Oper	The operational state of the service.
Mtu	The largest frame size (in octets) that the port can handle.
Adm	The largest frame size (in octets) that the SAP can handle.
SAP Count	The number of SAPs defined on the service.
SAP Type	The type of SAPs allowed in the service. It also describes the applied processing by the node to the packets received on these SAPs.
Identifier	Specifies the service access (SAP).
OprMTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this port, without requiring the packet to be fragmented.

Label	Description (Continued)
Opr	The operating state of the SAP

Sample Output

```
*A:7210SAS# show service id 10 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 10                  Vpn Id            : 0
Service Type      : VPLS
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 02/06/2106 06:28:12
Last Mgmt Change  : 01/10/1970 01:55:31
Admin State       : Down                Oper State        : Down
MTU               : Not Applicable      Def. Mesh VC Id   : 10
SAP Count         : 0
Uplink Type:      : L2
SAP Type:         : Dot1q Range         Customer vlan:    : n/a

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:7210SAS# show service id 10 base

A:Dut-A# show service id 1 base
===============================================
Service Basic Information
===============================================================================
Service Id : 1 Vpn Id : 0
Service Type : Epipe
Customer Id : 1
Last Status Change: 06/24/2001 00:57:55
Last Mgmt Change : 06/24/2001 00:51:36
Admin State : Up Oper State : Up
MTU : 1514
MTU Check : Disabled
Vc Switching : False
SAP count : 1 SDP Bind Count : 1
-------------------------------------------------------------------------------
Service Access and Destination Points
-------------------------------------------------------------------------------
Identifier Type    AdmMTU   OprMTU   Adm Opr
-------------------------------------------------------------------------------
sap:1/1/21:1 q-tag 1518     1518     Up Up
sdp:1:1 S<100.1.12> n/a 1518 1518 Up Up
=============================================================-----------------
A:Dut-A#
```

fdb

Syntax **fdb** [**sap** *sap-id* [**expiry**]] | [**mac** *ieee-address* [**expiry**]] | [**detail**] [**expiry**]

Context show>service>id
show>service>fdb-mac

Description This command displays FDB entries for a given MAC address.

Parameters **sap** *sap-id* — Specifies the physical port identifier portion of the SAP. See Common CLI Command Descriptions on page 987 for command syntax.

detail — Displays detailed information.

expiry — Displays time until MAC is aged out.

Show FDB Information — The following table describes service FDB output fields.

Label	Description
ServID	Displays the service ID.
MAC	Displays the associated MAC address.
Mac Move	Displays the administrative state of the MAC movement feature associated with this service.
Primary Factor	Displays a factor for the primary ports defining how many MAC-relearn periods should be used to measure the MAC-relearn rate.
Secondary Factor	Displays a factor for the secondary ports defining how many MAC-relearn periods should be used to measure the MAC-relearn rate.
Mac Move Rate	Displays the maximum rate at which MAC's can be re-learned in this 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 MAs. The rate is computed as the maximum number of re-learns allowed in a 5 second interval: for example, the default rate of 2 re-learns per second corresponds to 10 re-learns in a 5 second period.
Mac Move Timeout	Displays the time in seconds to wait before a SAP that has been disabled after exceeding the maximum 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.
Mac Move Retries	Displays the number of times retries are performed for reenabling the SAP/SDP.

Label	Description
Table Size	Specifies the maximum number of learned and static entries allowed in the FDB of this service.
Total Count	Displays the total number of learned entries in the FDB of this service.
Learned Count	Displays the current number of learned entries in the FDB of this service.
Static Count	Displays the current number of static entries in the FDB of this service.
OAM-learned Count	Displays the current number of OAM entries in the FDB of this service.
Remote Age	Displays 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	Displays the number of seconds used to age out FDB entries learned on local SAPs.
High Watermark	Displays the utilization of the FDB table of this service at which a table full alarm will be raised by the agent.
Low Watermark	Displays the utilization of the FDB table of this service at which a table full alarm will be cleared by the agent.
Mac Learning	Specifies whether the MAC learning process is enabled
Discard Unknown	Specifies whether frames received with an unknown destination MAC are discarded.
Mac Aging	Indicates whether the MAC aging process is enabled.
Relearn Only	Displays, that when enabled, either the FDB table of this service is full, or that the maximum system-wide number of MA's supported by the agent has been reached, and thus MAC learning is temporary disabled, and only MAC re-learns can take place.
Mac Subnet Len	Displays the number of bits to be considered when performing MAC-learning or MAC-switching.
Source-Identifier	The location where the MAC is defined.
Type/Age	Type – Specifies the number of seconds used to age out TLS FDB entries learned on local SAPs. Age – Specifies the number of seconds used to age out TLS FDB entrieslearned on an SDP. These entries correspond to MAC addresses learned on remote SAPs. L — Learned - Dynamic entries created by the learning process. OAM — Entries created by the OAM process.

Label	Description
	Static — Statically configured.
Last Change	Indicates the time of the most recent state changes.

Sample Output

```
A:Dut-A# show service id 305 fdb
===============================================================================
Forwarding Database, Service 305
===============================================================================
Service Id        : 305                Mac Move          : Disabled
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Table Size        : 500                Total Count       : 375
Learned Count     : 375                Static Count      : 0
Remote Age        : 60                 Local Age         : 60
High WaterMark    : 95%                Low Watermark     : 90%
Mac Learning      : Enabl              Discard Unknown   : Dsabl
Mac Aging         : Enabl              Relearn Only      : False
===============================================================================
A:Dut-A#
```

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. See Common CLI Command Descriptions on page 987 for command syntax.

summary — Displays summary host information.

labels

Syntax **labels**

Context show>service>id

Description This command 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.

Label	Description
Type	Indicates whether the SDP is spoke.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
E. Lbl	The VC label used by this device to send packets to the far-end device in this service by the SDP.

Sample Output

```
A:Dut-A# show service id 305 labels
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Binding        Type  I.Lbl                E.Lbl
-------------------------------------------------------------------------------
305        1217:305           Spok  130506               130516
305        1317:305           Spok  130454               130591
305        1417:305           Spok  130428               131015
305        1617:305           Spok  131060               130843
-------------------------------------------------------------------------------
Number of Bound SDPs : 4
===============================================================================
A:Dut-A#
```

l2pt

Syntax **l2pt disabled**
l2pt [**detail**]

Context show>service>id

Description This command displays Layer 2 Protocol Tunnel (L2-PT) route information associated with this service.

Parameters **disabled —** Displays only entries with termination disabled. This helps identify configuration errors.

detail — Displays detailed information.

Output **Show L2PT Fields —** The following table describes show L2PT output fields:

Label	Description
Service id	Displays the 24 bit (0..16777215) service instance identifier for the service.
L2pt-term enabled	Indicates if L2-PT-termination and/or Bpdu-translation is in use in this service by at least one SAP or spoke SDP binding. If in use, at least one of L2PT-termination or Bpdu-translation is enabled. When enabled it is not possible to enable STP on this service.

Label	Description (Continued)
L2pt-term disabled	Indicates that L2-PT-termination is disabled.
Bpdu-trans auto	Specifies the number of L2-PT PDU's are translated before being sent out on a port or sap.
Bpdu-trans disabled	Indicates that Bpdu-translation is disabled.
SAPs	Displays the number of SAPs with L2PT or BPDU translation enabled or disabled.
SDPs	Displays the number of SDPs with L2PT or BPDU translation enabled or disabled.
Total	Displays the column totals of L2PT entities.
SapId	The ID of the access point where this SAP is defined.
L2pt-termination	Indicates whether L2pt termination is enabled or disabled.
Admin Bpdu-translation	Specifies whether Bpdu translation is administratively enabled or disabled.
Oper Bpdu-translation	Specifies whether Bpdu translation is operationally enabled or disabled.
SdpId	Specifies the SAP ID.

Sample:

```
*A:7210SAS>show>service# id 1 l2pt detail

===============================================================================
L2pt details, Service id 1
===============================================================================

Service Access Points
-------------------------------------------------------------------------------
SapId                       L2pt-                       Admin Bpdu-  Oper Bpdu-
                            termination                 translation  translation
-------------------------------------------------------------------------------
1/1/1                       stp cdp vtp dtp pagp udld   disabled     disabled
-------------------------------------------------------------------------------
Number of SAPs : 1
===============================================================================

===============================================================================
L2pt summary, Service id 1
===============================================================================
       L2pt-term   L2pt-term   Bpdu-trans  Bpdu-trans  Bpdu-trans  Bpdu-trans
       enabled     disabled    auto        disabled    pvst        stp
-------------------------------------------------------------------------------
SAP's  1           0           0           1           0           0
SDP's  0           0           0           0           0           0
-------------------------------------------------------------------------------
```

```
Total  1           0           0           1           0           0
===============================================================================
*A:7210SAS>show>service#
```

mac-move

Syntax **mac-move**

Context show>service>id

Description This command displays MAC move related information about the service.

mac-protect

Syntax **mac-protect**

Context show>service>id

Description This command displays MAC protect-related information about the service.

mrouters

Syntax **mrouters** [**detail**]

Context show>service>id>mld-snooping

Description This command displays all multicast routers.

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.

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

Label	Description
Region Name	Displays the MSTP region name.
Region Revision	Displays the MSTP region revision.
MST Max Hops	Displays the MSTP maximum hops specified.

Label	Description (Continued)
Instance	Displays the MSTP instance number.
Priority	Displays the MSTP priority.
Vlans mapped	Displays the VLAN range of the MSTP instance.

Sample Output

```
*A:SASMX>show>service>id# mstp-configuration

============================================
Mstp configuration info, Service 5
============================================
Region Name        : abc
Region Revision    : 0
MST Max Hops       : 20

================================================================================
vlan to MST instance mapping
================================================================================
Instance  Priority  Vlans mapped
--------------------------------------------------------------------------------
2         0
================================================================================
*A:SASMX>show>service>id#
```

sap

Syntax **sap** *sap-id* [*filter*]

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. See Common CLI Command Descriptions on page 987 for command syntax.

detail — Displays detailed information for the SAP.

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 (Continued)
Admin State	The administrative state of the SAP.
Oper State	The operational state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, L2OperDown, RelearnLimitExceeded, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
Last Status Change	Specifies the time of the most recent operating status change to this SAP
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SAP.
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP.
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP.
Egress Filter-Id	The egress filter policy ID assigned to the SAP.
Acct. Pol	The accounting policy ID assigned to the SAP.
Collect Stats	Specifies whether collect stats is enabled.
SAP per Meter stats	
Ingress Meter	Specifies the meter ID.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded.
For. OutProf	The number of out-of-profile packets and octets. (rate above CIR and below PIR) forwarded by the ingress meter.
Ingress TD Profile	The profile ID applied to the ingress SAP.
Egress TD Profile	The profile ID applied to the egress SAP.
Alarm Cell Handling	The indication that OAM cells are being processed.
AAL-5 Encap	The AAL-5 encapsulation type.
Aggregate Policer	rate-indicates the rate of the aggregate policer. burst-indicates the burst-size of the aggregate policer.
Loopback Mode	Displays the Ethernet port loopback mode

Label	Description (Continued)
Loopback Src Addr	Displays the configured loopback source address
Loopback Dst Addr	Displays the configured loopback destination address
No-svc-port used	Displays the port ID of the port on which no service is configured. This port is used for the port loop back with MAC swap functionality.

Sample Output

```
A:7210>show>service>id# sap 1/1/1:1 detail

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/1/1:1                  Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
Last Status Change : 10/05/2010 07:22:04
Last Mgmt Change   : 10/05/2010 07:22:05
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100

Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
Admin MTU          : 1518                     Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
tod-suite          : None
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled

Acct. Pol          : None                     Collect Stats     : Disabled



-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                       Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : 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
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                        Min Auth Interval : 15 minutes

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 5                        Egress qos-policy : 1
-------------------------------------------------------------------------------
Aggregate Policer (Not Available)
-------------------------------------------------------------------------------
rate               : n/a                      burst             : n/a
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
Classifiers Allocated: 256                   Meters Allocated   : 32
Classifiers Used     : 2                     Meters Used        : 2
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                       Packets                 Octets
Ingress Stats:         0                       0
Egress Stats:          0                       0
-------------------------------------------------------------------------------
Sap per Meter stats
-------------------------------------------------------------------------------
                       Packets                 Octets

Ingress Meter 1 (Unicast)
For. InProf          : 0                       0
For. OutProf         : 0                       0

Ingress Meter 11 (Multipoint)
For. InProf          : 0                       0
For. OutProf         : 0                       0
===============================================================================
*A:SAS-M-A0-2>show>service>id# sap 1/1/1:10.* detail

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/1/1:10.*              Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Up
Flags              : None
Last Status Change : 04/29/2001 06:59:15
Last Mgmt Change   : 04/28/2001 03:09:30
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100

Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
Admin MTU          : 1522                    Oper MTU          : 1522
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : 1                       Egr Mac Fltr-Id   : n/a
tod-suite          : None
```

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

Acct. Pol          : None                    Collect Stats     : Disabled



-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                      Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Forwarding
Port Number        : 2048                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : N/A
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                     Designated Port   : N/A

Forward transitions: 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                       MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                       Min Auth Interval : 15 minutes

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1
-------------------------------------------------------------------------------
Aggregate Policer
-------------------------------------------------------------------------------
rate               : n/a                     burst             : n/a
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
Classifiers Allocated: 4                    Meters Allocated   : 2
Classifiers Used     : 2                    Meters Used        : 2
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                       Packets                 Octets
Ingress Stats:         142761481188            9707780720784
Egress Stats:          0                       0
Extra-Tag Drop Stats:  n/a                     n/a
-------------------------------------------------------------------------------
Sap per Meter stats
-------------------------------------------------------------------------------
                       Packets                 Octets

Ingress Meter 1 (Unicast)
For. InProf          : 17                      1162
```

```
For. OutProf          : 0                     0

Ingress Meter 11 (Multipoint)
For. InProf           : 61                    4148
For. OutProf          : 142761547917          9707785259394
===============================================================================
```

sdp

Note : SDP commands are not supported by 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

Syntax **sdp** [*sdp-id* | **far-end** *ip-addr*] [**detail**]

Context show>service>id

Description This command displays information for the SDPs associated with the service. If no optional parameters are specified, a summary of all associated SDPs is displayed.

Parameters *sdp-id* — 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 spoke.
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.

Label	Description (Continued)
Far End	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The 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.

Sample Output

```
A:Dut-A>show>service>id# sdp 1217:305
===============================================================================
Service Destination Point (Sdp Id : 1217:305)
===============================================================================
SdpId            Type IP address      Adm     Opr       I.Lbl       E.Lbl
-------------------------------------------------------------------------------
1217:305         Spok 10.20.1.2       Up      Up        130506      130516
-------------------------------------------------------------------------------
Number of SDPs : 1
===============================================================================
A:Dut-A>show>service>id# sdp 1217:305  detail


A:Dut-A>show>service>id#
===============================================================================
Service Destination Point (Sdp Id : 1217:305) Details
-------------------------------------------------------------------------------
 Sdp Id 1217:305  -(10.20.1.2)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id             : 1217:305                 Type             : Spoke
VC Type            : Ether                    VC Tag           : n/a
Admin Path MTU     : 0                        Oper Path MTU    : 9186
Far End            : 10.20.1.2                Delivery         : MPLS

Admin State        : Up                       Oper State       : Up
Acct. Pol          : None                     Collect Stats    : Disabled
Managed by Service : 300                      Prune State      : Not Pruned
Managed by Spoke   : 1217:300
Ingress Label      : 130506                   Egress Label     : 130516
```

```
Admin ControlWord  : Not Preferred           Oper ControlWord  : False
Last Status Change : 07/07/2009 18:49:40     Signaling         : TLDP
Last Mgmt Change   : 07/07/2009 14:39:14     Force Vlan-Vc     : Disabled
Last Mgmt Change   : 07/07/2009 14:39:14
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0

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

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

Statistics           :
I. Fwd. Pkts.      : 13601                   I. Fwd. Octs.     : 10676338
E. Fwd. Pkts.      : 83776987                E. Fwd. Octets    : 51589499116

Associated LSP LIST :
Lsp Name           : A_B_17
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 08h31m06s
-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Down                    Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Forwarding
Port Number        : 2049                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : N/A
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                     Designated Port Id: 0

Fwd Transitions    : 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
A:Dut-A>show>service>id#
```

split-horizon-group

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

Context show>service>id

Description This command displays service split horizon groups.

stp

Syntax **stp** [**detail**]

Context show>service>id

Description This command 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
`Bridge-id`	Specifies the MAC address used to identify this bridge in the network.
`Bridge fwd delay`	Specifies how fast a bridge changes its state when moving toward the forwarding state.
`Bridge Hello time`	Specifies the amount of time between the transmission of Configuration BPDUs.
`Bridge max age`	Specifies the maximum age of 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.

Label	Description (Continued)
Root path cost	Specifies the cost of the path to the root bridge as seen from this bridge.
Root forward delay	Specifies how fast the root changes its state when moving toward the forwarding state.
Root hello time	Specifies the amount of time between the transmission of configuration BPDUs.
Root max age	Specifies the maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded.
Root priority	This object specifies the priority of the bridge that is currently selected as root-bridge for the network.
Root port	Specifies the port number of the port which 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.
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.
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.

Sample Output

```
A:Dut-A>show>service>id# stp
===============================================================================
Stp info, Service 305
===============================================================================
Bridge Id          : 00:0d.00:20:ab:cd:00:01  Top. Change Count : 5
Root Bridge        : This Bridge              Stp Oper State    : Up
Primary Bridge     : N/A                      Topology Change   : Inactive
Mode               : Rstp                     Last Top. Change  : 0d 08:35:16
Vcp Active Prot.   : N/A
Root Port          : N/A                      External RPC      : 0
===============================================================================
Stp port info
===============================================================================
Sap/Sdp Id         Oper-     Port-      Port-       Port-  Oper-  Link-  Active
```

```
                   State     Role       State       Num    Edge   Type   Prot.
-------------------------------------------------------------------------------
1/1/16:305         Up        Designated Forward     2048   False  Pt-pt  Rstp
lag-4:305          Up        Designated Forward     2000   False  Pt-pt  Rstp
1217:305           Up        N/A        Forward     2049   N/A    Pt-pt  N/A
1317:305           Up        N/A        Forward     2050   N/A    Pt-pt  N/A
1417:305           Up        N/A        Forward     2051   N/A    Pt-pt  N/A
1617:305           Pruned    N/A        Discard     2052   N/A    Pt-pt  N/A
===============================================================================
A:Dut-A>show>service>id#


A:Dut-A>show>service>id# stp detail
===============================================================================
Spanning Tree Information
===============================================================================
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up                       Core Connectivity : Down
Stp Admin State    : Up                       Stp Oper State    : Up
Mode               : Rstp                     Vcp Active Prot.  : N/A

Bridge Id          : 00:0d.00:20:ab:cd:00:01  Bridge Instance Id: 13
Bridge Priority    : 0                        Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 0d 08:35:29              Bridge Max Age    : 20
Top. Change Count  : 5                        Bridge Fwd Delay  : 15
MST region revision: 0                        Bridge max hops   : 20
MST region name    :

Root Bridge        : This Bridge
Primary Bridge     : N/A

Root Path Cost     : 0                        Root Forward Delay: 15
Rcvd Hello Time    : 2                        Root Max Age      : 20
Root Priority      : 13                       Root Port         : N/A
-------------------------------------------------------------------------------
Spanning Tree Sap/Spoke SDP Specifics
-------------------------------------------------------------------------------
SAP Identifier     : 1/1/16:305               Stp Admin State   : Up
Port Role          : Designated               Port State        : Forwarding
Port Number        : 2048                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : PVST
Root Guard         : Disabled                 Active Protocol   : Rstp
Last BPDU from     : 80:04.00:0a:1b:2c:3d:4e
CIST Desig Bridge  : This Bridge              Designated Port   : 34816
Forward transitions: 5                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 29                       RST BPDUs tx      : 23488
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0

SAP Identifier     : lag-4:305                Stp Admin State   : Up
Port Role          : Designated               Port State        : Forwarding
Port Number        : 2000                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Rstp
```

```
Last BPDU from     : 80:04.00:0a:1b:2c:3d:4e
CIST Desig Bridge  : This Bridge              Designated Port   : 34768
Forward transitions: 4                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 23                       RST BPDUs tx      : 23454
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0

SDP Identifier     : 1217:305                 Stp Admin State   : Down
Port Role          : N/A                      Port State        : Forwarding
Port Number        : 2049                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last 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 Identifier     : 1317:305                 Stp Admin State   : Down
Port Role          : N/A                      Port State        : Forwarding
Port Number        : 2050                     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 Identifier     : 1417:305                 Stp Admin State   : Down
Port Role          : N/A                      Port State        : Forwarding
Port Number        : 2051                     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    : 1                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0

SDP Identifier     : 1617:305                 Stp Admin State   : Down
Port Role          : N/A                      Port State        : Discarding
Port Number        : 2052                     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
===============================================================================
A:Dut-A>show>service>id#

*7210-SAS>show>service>id# stp detail

===============================================================================
Spanning Tree Information
===============================================================================

-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up                       Core Connectivity : Down
Stp Admin State    : Up                       Stp Oper State    : Up
Mode               : Mstp                     Vcp Active Prot.  : N/A

Bridge Id          : 80:00.00:25:ba:04:66:a0  Bridge Instance Id: 0
Bridge Priority    : 32768                    Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 0d 02:54:16              Bridge Max Age    : 20
Top. Change Count  : 27                       Bridge Fwd Delay  : 15

Root Bridge        : 40:00.7c:20:64:ac:ff:63
Primary Bridge     : N/A

Root Path Cost     : 10                       Root Forward Delay: 15
Rcvd Hello Time    : 2                        Root Max Age      : 20
Root Priority      : 16384                    Root Port         : 2048

MSTP info for CIST :
Regional Root      : 80:00.7c:20:64:ad:04:5f  Root Port         : 2048
Internal RPC       : 10                       Remaining Hopcount: 19
MSTP info for MSTI 1 :
Regional Root      : This Bridge              Root Port         : N/A
Internal RPC       : 0                        Remaining Hopcount: 20
MSTP info for MSTI 2 :
Regional Root      : 00:02.7c:20:64:ad:04:5f  Root Port         : 2048
Internal RPC       : 10                       Remaining Hopcount: 19


-------------------------------------------------------------------------------
Spanning Tree Sap Specifics
-------------------------------------------------------------------------------
SAP Identifier     : 1/1/7:0                  Stp Admin State   : Up
Port Role          : Root                     Port State        : Forwarding
Port Number        : 2048                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Mstp
Last BPDU from     : 80:00.7c:20:64:ad:04:5f  Inside Mst Region : True
CIST Desig Bridge  : 80:00.7c:20:64:ad:04:5f  Designated Port   : 34816
MSTI 1 Port Prio   : 128                      Port Path Cost    : 10
MSTI 1 Desig Brid  : This Bridge              Designated Port   : 34816
MSTI 2 Port Prio   : 128                      Port Path Cost    : 10
MSTI 2 Desig Brid  : 00:02.7c:20:64:ad:04:5f  Designated Port   : 34816
Forward transitions: 17                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
```

```
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 7310                    MST BPDUs tx      : 7277

SAP Identifier     : 1/1/8:0                 Stp Admin State   : Up
Port Role          : Alternate               Port State        : Discarding
Port Number        : 2049                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : False
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : Mstp
Last BPDU from     : 80:00.7c:20:64:ad:04:5f Inside Mst Region : True
CIST Desig Bridge  : 80:00.7c:20:64:ad:04:5f Designated Port   : 34817
MSTI 1 Port Prio   : 128                     Port Path Cost    : 10
MSTI 1 Desig Brid  : This Bridge             Designated Port   : 34817
MSTI 2 Port Prio   : 128                     Port Path Cost    : 10
MSTI 2 Desig Brid  : 00:02.7c:20:64:ad:04:5f Designated Port   : 34817
Forward transitions: 14                      Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 7326                    MST BPDUs tx      : 7307

SAP Identifier     : 1/1/9:0                 Stp Admin State   : Up
Port Role          : Designated              Port State        : Forwarding
Port Number        : 2050                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : True
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : Mstp
Last BPDU from     : N/A                     Inside Mst Region : True
CIST Desig Bridge  : This Bridge             Designated Port   : 34818
MSTI 1 Port Prio   : 128                     Port Path Cost    : 10
MSTI 1 Desig Brid  : This Bridge             Designated Port   : 34818
MSTI 2 Port Prio   : 128                     Port Path Cost    : 10
MSTI 2 Desig Brid  : This Bridge             Designated Port   : 34818
Forward transitions: 2                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                       MST BPDUs tx      : 7415

SAP Identifier     : 1/1/25:0                Stp Admin State   : Up
Port Role          : Alternate               Port State        : Discarding
Port Number        : 2051                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : False
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : Mstp
Last BPDU from     : 80:00.7c:20:64:ad:04:5f Inside Mst Region : True
CIST Desig Bridge  : 80:00.7c:20:64:ad:04:5f Designated Port   : 34820
MSTI 1 Port Prio   : 128                     Port Path Cost    : 10
MSTI 1 Desig Brid  : This Bridge             Designated Port   : 34819
MSTI 2 Port Prio   : 128                     Port Path Cost    : 10
MSTI 2 Desig Brid  : 00:02.7c:20:64:ad:04:5f Designated Port   : 34820
Forward transitions: 10                      Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 7329                    MST BPDUs tx      : 7303

SAP Identifier     : lag-1:0                 Stp Admin State   : Up
```

```
Port Role          : Alternate               Port State        : Discarding
Port Number        : 2052                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : False
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : Mstp
Last BPDU from     : 80:00.7c:20:64:ad:04:5f Inside Mst Region : True
CIST Desig Bridge  : 80:00.7c:20:64:ad:04:5f Designated Port   : 34822
MSTI 1 Port Prio   : 128                     Port Path Cost    : 10
MSTI 1 Desig Brid  : This Bridge             Designated Port   : 34820
MSTI 2 Port Prio   : 128                     Port Path Cost    : 10
MSTI 2 Desig Brid  : 00:02.7c:20:64:ad:04:5f Designated Port   : 34822
Forward transitions: 11                      Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 7322                    MST BPDUs tx      : 7299

===============================================================================
```

Sample Output

```
Sample output with MSTP information for 7210 SAS-M:

*A:SASMX[S0]>show>service>id# stp mst-instance 2

===============================================================================
MSTP specific info for service 5 MSTI 2
===============================================================================
Regional Root      : N/A                     Root Port         : N/A
Internal RPC       : 0                       Remaining Hopcount: 20

===================================================================
MSTP port info for MSTI 2
===================================================================
Sap/Sdp Id         Oper-     Port-      Port-       Port-  Same
                   State     Role       State       Num    Region
-------------------------------------------------------------------
No data found.
===================================================================
*A:SASMX[S0]>show>service>id#

Sample output with MSTP information for 7210 SAS-M:

*A:SASMX[S0]>show>service>id# stp mst-instance 2

===============================================================================
MSTP specific info for service 5 MSTI 2
===============================================================================
Regional Root      : N/A                     Root Port         : N/A
Internal RPC       : 0                       Remaining Hopcount: 20

===================================================================
MSTP port info for MSTI 2
===================================================================
Sap/Sdp Id         Oper-     Port-      Port-       Port-  Same
                   State     Role       State       Num    Region
-------------------------------------------------------------------
```

```
No data found.
===============================================================================
*A:SASMX[S0]>show>service>id#
```

dhcp

Syntax **dhcp**

Context show>service>id

Description This command enables the context to display DHCP information for the specified service.

statistics

Syntax **statistics** [**sap** *sap-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.

interface *interface-name* — Displays information 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. Includes DHCP packets received from both DHCP client and DHCP server.
Transmitted Packets	The number of packets transmitted to the DHCP clients. Includes DHCP packets transmitted from both DHCP client and DHCP server.
Received Malformed Packets	The number of corrupted/invalid packets received from the DHCP clients. Includes DHCP packets received from both DHCP client and DHCP server
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients. In this case, a frame is dropped 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.

Label	Description
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.

```
*A:7210SAS>show>service>id>dhcp# statistics

===============================================================================
DHCP Global Statistics, service 1
===============================================================================
Rx Packets                          : 416554
Tx Packets                          : 206405
Rx Malformed Packets                : 0
Rx Untrusted Packets                : 0
Client Packets Discarded            : 0
Client Packets Relayed              : 221099
Client Packets Snooped              : 0
Client Packets Proxied (RADIUS)     : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded            : 0
Server Packets Relayed              : 195455
Server Packets Snooped              : 0
DHCP RELEASEs Spoofed               : 0
DHCP FORCERENEWs Spoofed            : 0
===============================================================================
*A:7210SAS>show>service>id>dhcp#
```

summary

Syntax **summary [interface** *interface-name]*

Context show>service>id>dhcp

Description Displays DHCP configuration summary information.

Parameters *interface interface-name* — Displays information for the specified IP 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	Specifies whether or not ARP populate is enabled. 7210 SAS does not support ARP populate.
Used/Provided	7210 SAS does not maintain lease state.
Info Option	Indicates whether Option 82 processing is enabled on the interface.
Admin State	Indicates the administrative state.

Sample Output

```
A:7210SAS# show service id 1 dhcp summary
DHCP Summary, service 1
===============================================================================

Interface Name                   Arp      Used/                Info    Admin
  SapId/Sdp                      Populate Provided             Option  State
-------------------------------------------------------------------------------
egr_1                            No       0/0                  Replace Up
i_1                              No       0/0                  Replace Up
-------------------------------------------------------------------------------
Interfaces: 2
===============================================================================

*A:7210SAS>show>service>id>dhcp#
```

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 This command 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 or SDP Id	Displays the SAP or SDP IDs of the service ID.
Oper State	Displays the operational state of the SAP or 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 MVR Groups	Specifies the actual number of multicast groups that can be joined on this SAP or SDP.
MVR From VPLS Cfg Drops	Displays the from VPLS drop count.
MVR To SAP Cfg Drops	Displays the to SAP drop count.
MVR Admin State	Displays the administrative state of MVR.

Label	Description (Continued)
MVR Policy	The MVR policy name.

Sample Output

```
*Sample output (7210 SAS-M and 7210 SAS-T in network mode)

*A:7210-SAS>show>service>id>igmp-snooping# all

===============================================================================
IGMP Snooping info for service 2
===============================================================================

-------------------------------------------------------------------------------
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp                Oper    MRtr Send    Max  MVR       Num
Id                     State   Port Queries Grps From-VPLS Grps
-------------------------------------------------------------------------------
sap:1/1/1              Up      No   No      None 1         1
sap:1/1/4              Up      No   No      None Local     0

-------------------------------------------------------------------------------
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    Up Time
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of groups: 0

-------------------------------------------------------------------------------
IGMP Snooping SAP 1/1/1 Port-DB
-------------------------------------------------------------------------------
Group Address   Type    From-VPLS  Up Time        Expires  MC
                                                          Stdby
-------------------------------------------------------------------------------
224.1.1.1       dynamic 1          0d 00:11:01    246s
-------------------------------------------------------------------------------
Number of groups: 1

-------------------------------------------------------------------------------
```

```
IGMP Snooping SAP 1/1/4 Port-DB
-------------------------------------------------------------------------------
Group Address   Type    From-VPLS  Up Time        Expires  MC
                                                          Stdby
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of groups: 0

-------------------------------------------------------------------------------
IGMP Snooping Static Groups
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
IGMP Snooping Statistics
-------------------------------------------------------------------------------
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         0             0             0
Group Queries           0             0             0
V1 Reports              0             0             0
V2 Reports              68165         0             0
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length               : 0
Bad IP Checksum          : 0
Bad IGMP Checksum        : 0
Bad Encoding             : 0
No Router Alert          : 0
Zero Source IP           : 0
Wrong Version            : 0
Lcl-Scope Packets        : 0

Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
MCS Failures             : 0

MVR From VPLS Cfg Drops  : 68129
MVR To SAP Cfg Drops     : 0

-------------------------------------------------------------------------------
IGMP Snooping Multicast VPLS Registration info
-------------------------------------------------------------------------------
IGMP Snooping Admin State : Down

MVR Admin State           : Down
MVR Policy                : None
-------------------------------------------------------------------------------
Local SAPs/SDPs
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp               Oper      From       Num Local
           Id                    State     VPLS       Groups
-------------------------------------------------------------------------------
2          sap:1/1/1             Up        1          0
2          sap:1/1/4             Up        Local      0
-------------------------------------------------------------------------------
MVR SAPs (from-vpls=2)
-------------------------------------------------------------------------------
```

```
Svc Id     Sap/Sdp                Oper      From       Num MVR
           Id                     State     VPLS       Groups
-------------------------------------------------------------------------------
No MVR SAPs found.
===============================================================================
*A:7210-SAS>show>service>id>igmp-snooping#
```

Sample output (7210 SAS-M in access-uplink mode):

```
A:7210-SAS>show>service>id# igmp-snooping all

===============================================================================
IGMP Snooping info for service 1
===============================================================================

-------------------------------------------------------------------------------
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Up
Querier     : 1.1.1.1 on SAP 1/1/1
-------------------------------------------------------------------------------
Sap/Sdp                Oper    MRtr Send    Max  Max  Num
Id                     State   Port Queries Grps Srcs Grps
-------------------------------------------------------------------------------
sap:1/1/1              Up      Yes  No      None None 0
sap:1/1/2              Up      No   No      None None 1

-------------------------------------------------------------------------------
IGMP Snooping Querier info
-------------------------------------------------------------------------------
Sap Id                  : 1/1/1
IP Address              : 1.1.1.1
Expires                 : 255s
Up Time                 : 0d 16:51:04
Version                 : 2

General Query Interval  : 125s
Query Response Interval : 10.0s
Robust Count            : 2

-------------------------------------------------------------------------------
IGMP Snooping Multicast Routers
-------------------------------------------------------------------------------
MRouter          Sap/Sdp Id             Up Time        Expires   Version
-------------------------------------------------------------------------------
1.1.1.1          1/1/1                  0d 16:51:14    255s      2
-------------------------------------------------------------------------------
Number of mrouters: 1

-------------------------------------------------------------------------------
IGMP Snooping Proxy-reporting DB
-------------------------------------------------------------------------------
Group Address    Mode     Up Time        Num Sources
-------------------------------------------------------------------------------
224.1.1.2        exclude  0d 16:51:14    0
-------------------------------------------------------------------------------
Number of groups: 1
```

```
-------------------------------------------------------------------------------
IGMP Snooping SAP 1/1/1 Port-DB
-------------------------------------------------------------------------------
Group Address   Mode    Type    Up Time        Expires  Num
                                                        Src
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of groups: 0

-------------------------------------------------------------------------------
IGMP Snooping SAP 1/1/2 Port-DB
-------------------------------------------------------------------------------
Group Address   Mode    Type    Up Time        Expires  Num
                                                        Src
-------------------------------------------------------------------------------
224.1.1.2       exclude dynamic 0d 16:51:17    259s     0
-------------------------------------------------------------------------------
Number of groups: 1

-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
IGMP Snooping Statistics
-------------------------------------------------------------------------------
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         811311        0             811311
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              18030         11928         0
V3 Reports              0             0             0
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length               : 0
Bad IP Checksum          : 0
Bad IGMP Checksum        : 0
Bad Encoding             : 0
No Router Alert          : 0
Zero Source IP           : 0
Wrong Version            : 0
Lcl-Scope Packets        : 0

Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
Exceeded Max Num Sources : 0
===============================================================================
```

mfib

Syntax **mfib** [**brief**] [**ip** | **mac**] **brief**
mfib [**group** *grp-address*]

Context show>service>id

Description This command displays the multicast FIB on the VPLS service.

Parameters **brief —** Displays a brief output.

group grp *grp-address* **—** Displays the multicast FIB for a specific multicast group address.

Output **Show Output —** The following table describes the command output fields:

Label	Description
Group Address	IPv4 multicast group address.
SAP ID	Indicates the SAP/SDP to which the corresponding multicast stream will be forwarded/blocked.
Forwarding/Blocking	Indicates whether the corresponding multicast stream will be blocked/forwarded.
Number of Entries	Specifies the number of entries in the MFIB.
Forwarded Packets	Indicates the number of multicast packets forwarded for the corresponding source/group.
Forwarded Octets	Indicates the number of octets forwarded for the corresponding source/group.
Svc ID	Indicates the service to which the corresponding multicast stream will forwarded/blocked. Local means that the multicast stream will be forwarded/blocked to a SAP or SDP local to the service.

Sample Output

```
*A:SAS# show service id 1 mfib

===============================================================================
Multicast FIB, Service 1
===============================================================================
Group Address         Sap/Sdp Id               Svc Id   Fwd/Blk
-------------------------------------------------------------------------------
224.4.4.4             sap:1/1/1                Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
A:7210-SAS>show>service>id#
```

mrouters

Syntax **mrouters** [**detail**]

Context show>service>id>igmp-snooping

Description This command displays all multicast routers.

Parameters **detail —** Displays detailed information.

Sample Output

```
A:7210-SAS>show>service>id# igmp-snooping mrouters

===============================================================================
IGMP Snooping Multicast Routers for service 1
===============================================================================
MRouter          Sap/Sdp Id             Up Time        Expires   Version
-------------------------------------------------------------------------------
1.1.1.1          1/1/1                  0d 16:53:44    254s      2
-------------------------------------------------------------------------------
Number of mrouters: 1
===============================================================================
A:7210-SAS>show>service>id#
```

mvr

Syntax **mvr**

Context show>service>id>igmp-snooping

Description This command displays Multicast VPLS Registration (MVR) information.

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 svcid.
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.

Label	Description
Num Groups	Specifies the number of groups learned via this local SAP.

Sample output

```
*A:7210-SAS>show>service>id>igmp-snooping# mvr

===============================================================================
IGMP Snooping Multicast VPLS Registration info for service 2
===============================================================================
IGMP Snooping Admin State : Down

MVR Admin State           : Down
MVR Policy                : None
-------------------------------------------------------------------------------
Local SAPs/SDPs
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp               Oper      From       Num Local
           Id                    State     VPLS       Groups
-------------------------------------------------------------------------------
2          sap:1/1/1             Up        1          0
2          sap:1/1/4             Up        Local      0
-------------------------------------------------------------------------------
MVR SAPs (from-vpls=2)
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp               Oper      From       Num MVR
           Id                    State     VPLS       Groups
-------------------------------------------------------------------------------
No MVR SAPs found.
===============================================================================
*A:7210-SAS>show>service>id>igmp-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 This command 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. See Common CLI Command Descriptions on page 987 for command syntax.

sdp *sdp-id* **—** Displays only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 — 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

Output **Show Output —** The following table describes the show output fields:

Label	Description
Group Address	The IP multicast group address for which this entry contains information.
Mode	Specifies 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.
Type	Indicates how this group entry was learned. If this group entry was learned by IGMP, the value is set to dynamic. For statically configured groups, the value is set to static.
Compatibility mode	Specifies the IGMP mode. This is 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.
V1 host expires	The time remaining until the local router will assume that there are no longer any IGMP Version 1 members on the IP subnet attached to this interface. Upon hearing any IGMPv1 membership report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv2 leave messages for this group that it receives on this interface.

Label	Description
V2 host expires	The time remaining until the local router will assume that there are no longer any IGMP Version 2 members on the IP subnet attached to this interface. Upon hearing any IGMPv2 membership report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv3 leave messages for this group that it receives on this interface.
Source address	The source address for which this entry contains information.
Up Time	The time since the source group entry was created.
Expires	The amount of time remaining before this entry will be aged out.
Number of sources	Indicates the number of IGMP group and source specific queries received on this SAP.
Forwarding/Blocking	Indicates whether this entry is on the forward list or block list.
Number of groups	Indicates the number of groups configured for this SAP.
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.

Sample Output (for 7210 SAS-M and 7210 SAS-T devices configured in network mode)

```
*A:7210-SAS>show>service>id>igmp-snooping# port-db sap 1/1/1

===============================================================================
IGMP Snooping SAP 1/1/1 Port-DB for service 2
===============================================================================
Group Address   Type    From-VPLS  Up Time        Expires  MC
                                                          Stdby
-------------------------------------------------------------------------------
224.1.1.1       dynamic 1          0d 00:15:57    246s
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
*A:7210-SAS>show>service>id>igmp-snooping#
===============================================================
*A:MTU-7210#
*A:7210-SAS>show>service>id>igmp-snooping# port-db sap 1/1/1 detail

===============================================================================
IGMP Snooping SAP 1/1/1 Port-DB for service 2
===============================================================================
-------------------------------------------------------------------------------
IGMP Group 224.1.1.1
-------------------------------------------------------------------------------
Type             : dynamic
Up Time          : 0d 00:14:30          Expires          : 259s
Compat Mode      : IGMP Version 2
V1 Host Expires  : 0s                   V2 Host Expires  : 259s
MVR From-VPLS    : 1                    MVR To-SAP       : 1/1/4
```

```
MC Standby       : no
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
*A:7210-SAS>show>service>id>i
```

proxy-db

Syntax **proxy-db** [**detail**]
proxy-db group *grp-address*

Context show>service>id>igmp-snooping

Description This command 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.

Output **Show Output —** The following table describes the show output fields:

Label	Description
Group Address	The IP multicast group address for which this entry contains information.
Mode	Specifies 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 the “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.
Up Time	The total operational time in seconds.
Number of groups	Number of IGMP groups.

Sample Output

```
*A:MTU-7210#  show service id 100 igmp-snooping proxy-db
===============================================================================
IGMP Snooping Proxy-reporting DB for service 100
===============================================================================
Group Address    Up Time
-------------------------------------------------------------------------------
227.7.7.7        0d 00:05:30
227.7.7.8        0d 00:05:30
228.8.8.8        0d 00:03:42
-------------------------------------------------------------------------------
Number of groups: 3
```

```
===============================================================================
*A:MTU-7210#


*A:MTU-T2#  show service id 100 igmp-snooping proxy-db detail
===============================================================================
IGMP Snooping Proxy-reporting DB for service 100
===============================================================================
-------------------------------------------------------------------------------
IGMP Group 227.7.7.7
-------------------------------------------------------------------------------
Up Time : 0d 00:05:43
-------------------------------------------------------------------------------
IGMP Group 227.7.7.8
-------------------------------------------------------------------------------
Up Time : 0d 00:05:43
-------------------------------------------------------------------------------
IGMP Group 228.8.8.8
-------------------------------------------------------------------------------
Up Time : 0d 00:03:55
-------------------------------------------------------------------------------
Number of groups: 3
===============================================================================
*A:MTU-7210#
```

querier

Syntax **querier**

Context show>service>id>igmp-snooping

Description This command displays information on the IGMP snooping queriers for the VPLS service.

Output **Show Output —** The following table describes the show output fields:

Label	Description
SAP Id	Specifies the SAP ID of the service.
IP address	Specifies the IP address of the querier.
Expires	The time left, in seconds, that the query will expire.
Up time	The length of time the query has been enabled.
Version	The configured version of IGMP.
General Query Interval	The frequency at which host-query packets are transmitted.
Query Response Interval	The time to wait to receive a response to the host-query message from the host.
Robust Count	Specifies the value used to calculate several IGMP message intervals.

Label	Description (Continued)

Sample Output

```
*A:MTU-7210#  show service id 100 igmp-snooping querier
===============================================================================
IGMP Snooping Querier info for service 100
===============================================================================
Sap Id                  : 1/1/1
IP Address              : 10.10.9.9
Expires                 : 24s
Up Time                 : 0d 00:05:20
Version                 : 2

General Query Interval  : 10s
Query Response Interval : 10.0s
Robust Count            : 2
===============================================================================
*A:MTU-7210#


*A:MTU-T2#  show service id 100 igmp-snooping proxy-db
===============================================================================
IGMP Snooping Proxy-reporting DB for service 100
===============================================================================
Group Address    Up Time
-------------------------------------------------------------------------------
227.7.7.7        0d 00:05:30
227.7.7.8        0d 00:05:30
228.8.8.8        0d 00:03:42
-------------------------------------------------------------------------------
Number of groups: 3
===============================================================================
*A:MTU-T2#
```

static

Syntax **static** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>igmp-snooping

Description This command 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. See Common CLI Command Descriptions on page 987 for command syntax.

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

Output **Show Output —** The following table describes the show output fields:

Label	Description
Source	Displays the IP source address used in IGMP queries.
Group	Displays the static IGMP snooping source groups for a specified SAP.

Sample Output

```
*A:MTU-7210#  show service id 100 igmp-snooping static

===============================================================================
IGMP Snooping Static Groups for service 100
===============================================================================
-------------------------------------------------------------------------------
IGMP Snooping Static Groups for SAP 1/1/2
-------------------------------------------------------------------------------
Group
-------------------------------------------------------------------------------
228.8.8.8
-------------------------------------------------------------------------------
Static (*,G) entries: 1
===============================================================================
*A:MTU-7210#
```

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>igmp-snooping

Description This command displays IGMP snooping statistics for the VPLS service.

Parameters **sap** *sap-id* **—** Displays IGMP snooping statistics for a specific SAP. See Common CLI Command Descriptions on page 987 for command syntax.

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

Sample Output

Sample Output (SAS-M in network mode)

```
*A:7210-SAS>show>service>id>igmp-snooping# statistics

===============================================================================
IGMP Snooping Statistics for service 2
===============================================================================
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         0             0             0
Group Queries           0             0             0
V1 Reports              0             0             0
V2 Reports              142207        0             0
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length               : 0
Bad IP Checksum          : 0
Bad IGMP Checksum        : 0
Bad Encoding             : 0
No Router Alert          : 0
Zero Source IP           : 0
Wrong Version            : 0
Lcl-Scope Packets        : 0

Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
MCS Failures             : 0

MVR From VPLS Cfg Drops  : 142130
MVR To SAP Cfg Drops     : 0
===============================================================================
*A:7210-SAS>show>service>id>igmp-snooping#
```

Sample Output (SAS-M and 7210 SAS-T in access-uplink mode)

```
A:7210-SAS>show>service>id# igmp-snooping statistics

===============================================================================
IGMP Snooping Statistics for service 1
===============================================================================
```

```
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         816014        0             816014
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              18134         11991         0
V3 Reports              0             0             0
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length               : 0
Bad IP Checksum          : 0
Bad IGMP Checksum        : 0
Bad Encoding             : 0
No Router Alert          : 0
Zero Source IP           : 0
Wrong Version            : 0
Lcl-Scope Packets        : 0

Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
Exceeded Max Num Sources : 0
===============================================================================
A:7210-SAS>show>service>id#
```

endpoint

Syntax **endpoint** [*endpoint-name*]

Context show>service>id

Description This command displays service endpoint information.

Parameters *endpoint-name* — Specifies an endpoint name created in the **config>service>vpls** context.

Sample Output

```
*A:Dut-B# show service id 1 endpoint
===============================================================================
Service 1 endpoints
===============================================================================
Endpoint name               : mcep-t1
Description                 : (Not Specified)
Revert time                 : 0
Act Hold Delay              : 0
Ignore Standby Signaling    : false
Suppress Standby Signaling  : false
Block On Mesh Fail          : true
Psv Mode Active             : No
Tx Active                   : 231:1
Tx Active Up Time           : 0d 00:06:57
Revert Time Count Down      : N/A
Tx Active Change Count      : 5
```

```
Last Tx Active Change        : 02/13/2009 22:08:33
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
Spoke-sdp: 221:1 Prec:1                             Oper Status: Up
Spoke-sdp: 231:1 Prec:2                             Oper Status: Up
===============================================================================
*A:Dut-B#
```

VPLS 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.

Values service-id: 1 — 214748364
svc-name: A string up to 64 characters in length.

statistics

Syntax **statistics**

Context clear>service>stats

Description This command clears session statistics for this service.

fdb

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

Context clear>service>id

Description This command clears FDB entries for the service.

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

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

sap-id — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 987 for command syntax.

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

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

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

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

Note : SDP commands are not supported by 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

Syntax **mesh-sdp** *sdp-id*[**:***vc-id*] **ingress-vc-label**

Context clear>service>id

Description This command 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

Note : SDP commands are not supported by 7210 SAS-M and 7210 SAS-T devices configured in Access uplink mode.

Syntax **spoke-sdp** *sdp-id*[:*vc-id*] {**all | counters | stp | l2pt**}}

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

all — Clears all queue statistics and STP statistics associated with the SDP.

counters — Clears all queue statistics associated with the SDP.

stp — Clears all STP statistics associated with the SDP.

l2pt — Clears all L2PT statistics associated with the SDP.

sap

Syntax **sap** *sap-id*

Context clear>service>statistics

Description This command clears statistics for the SAP bound to the service.

Parameters *sap-id* — See Common CLI Command Descriptions on page 987 for command syntax.

all — Clears all queue statistics and STP statistics associated with the SAP.

counters — Clears all queue statistics associated with the SAP.

counters

Syntax **counters**

Context clear>service>statistics>id

Description This command clears all traffic queue counters associated with the service ID.

l2pt

Syntax **l2pt**

Context clear>service>statistics>id

Description This command clears the l2pt statistics for this service.

mesh-sdp

Syntax **mesh-sdp** *sdp-id[:vc-id]* {**all | counters | stp | mrp**}

Context clear>service>statistics>id

Description This command clears the statistics for a particular mesh SDP bind.

Parameters *sdp-id[:vc-id]* — sdp-id - [1..17407]
vc-id - [1..4294967295]

all — Clears all queue statistics and STP statistics associated with the SDP.

counters — Clears all queue statistics associated with the SDP.

stp — Clears all STP statistics associated with the SDP.

mrp — Clears all MRP statistics associated with the SDP.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] {**all** | **counters** | **stp| l2pt**}

Context clear>service>statistics>id

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.

l2pt — Clears all L2PT 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*}

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.

Parameters **all —** Clears all detected protocol statistics.

sap-id — Clears the specified lease state SAP information. See Common CLI Command Descriptions on page 987 for command syntax.

igmp-snooping

Syntax **igmp-snooping**

Context clear>service>id

Description This command enables the context to clear IGMP snooping data.

port-db

Syntax **port-db** [**sap** *sap-id*] [**group** *grp-address*]
port-db sdp *sdp-id:vc-id* [**group** *grp-address*]

Context clear>service>id>igmp-snooping

Description This command 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. See Common CLI Command Descriptions on page 987 for command syntax.

sdp-id — Clears only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 — 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.

querier

Syntax **querier**

Context clear>service>id>igmp-snooping

Description This command clears the information on the IGMP snooping queriers for the VPLS service.

VPLS 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.

Values service-id: 1 — 214748364
svc-name: A string up to 64 characters in length.

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 a particular debugging event type.
The **no** form of the command disables the event type debugging.

Parameters **config-change —** Debugs configuration change events.

svc-oper-status-change — Debugs service operational status changes.

sap-oper-status-change — Debugs SAP operational status changes.

sdpbind-oper-status-change — Debugs SDP operational status changes.

sap

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

Context debug>service>id

Description This command enables debugging for a particular SAP.

Parameters *sap-id* — Specifies the SAP ID.

stp

Syntax **stp**

Context debug>service>id

Description This command enables the context for debugging STP.

all-events

Syntax **all-events**

Context debug>service>id>stp

Description This command enables STP debugging for all events.

bpdu

Syntax [**no**] **bpdu**

Context debug>service>id>stp

Description This command enables STP debugging for received and transmitted BPDUs.

core-connectivity

Syntax [**no**] **core-connectivity**

Context debug>service>id>stp

Description This command enables STP debugging for core connectivity.

exception

Syntax [**no**] **exception**

Context debug>service>id>stp

Description This command enables STP debugging for exceptions.

fsm-state-changes

Syntax [**no**] **fsm-state-changes**

Context debug>service>id>stp

Description This command enables STP debugging for FSM state changes.

fsm-timers

Syntax [**no**] **fsm-timers**

Context debug>service>id>stp

Description This command enables STP debugging for FSM timer changes.

port-role

Syntax [**no**] **port-role**

Context debug>service>id>stp

Description This command enables STP debugging for changes in port roles.

port-state

Syntax [**no**] **port-state**

Context debug>service>id>stp

Description This command enables STP debugging for port states.

sap

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

Context debug>service>id>stp

Description This command enables STP debugging for a specific SAP.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 987 for command syntax.

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>stp

Description This command enables STP debugging for a specific SDP.

Common CLI Command Descriptions

In This Chapter

This section provides information about common Command Line Interface (CLI) syntax and command usage.

Topics in this chapter include:

Common Service Commands

sap

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

Description This command specifies the physical port identifier portion of the SAP definition.

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

The *sap-id* can be configured in one of the following formats:

Type	Syntax	Example
port-id	*slot*/*mda*/*port*[.*channel*]	1/1/5
null	[*port-id* | *lag-id*]	*port-id*: 1/1/3 *lag-id*: lag-3
dot1q	[*port-id* | *lag-id*]:qtag1	*port-id*:qtag1: 1/1/3:100 *lag-id*:qtag1:lag-3:102 *cp*.conn-prof-id: 1/2/1:cp.2
qinq	[*port-id* | *lag-id*]:*qtag1.qtag2*	*port-id*:qtag1.qtag2: 1/1/3:100.10 *lag-id*:qtag1.qtag2: lag-10:

The values depends on the encapsulation type configured for the interface. The following table describes the allowed values for the port and encapsulation types.

Port Type	Encap-Type	Allowed Values	Comments
Ethernet	Null	0	The SAP is identified by the port.
Ethernet	Dot1q	0 — 4094	The SAP is identified by the 802.1Q tag on the port. Note that a 0 qtag1 value also accepts untagged packets on the dot1q port.
Ethernet	QinQ	qtag1: 0 — 4094 qtag2: 0 — 4094	The SAP is identified by two 802.1Q tags on the port. Note that a 0 qtag1 value also accepts untagged packets on the Dot1q port.

Appendix: Port-Based Split Horizon

In This Chapter

This section provides Port-Based Split Horizon configuration information.

Overview

The port-based split horizon feature can be used to disable local switching on the 7210 SAS. A loop-free topology can be achieved using split horizon on 7210 SAS switches.

Traffic arriving on an access or a network port within a split horizon group will not be copied to other access and a network ports in the same split horizon group, but will be copied to an access or network ports in other split horizon groups.

Since split horizon is a per port feature in 7210 SAS, all SAPs associated with the port becomes part of split horizon group configured on that port.

Topology

Figure illustrates an example of split horizon groups used to prevent communication between two access SAPs and between two network ports.

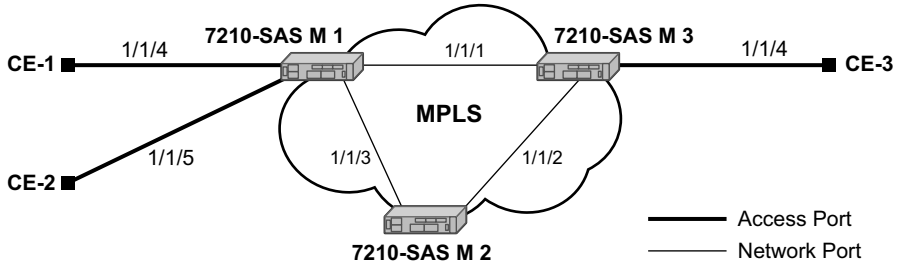

Figure 70: Split Horizon Group Example

Using 7210-SAS-1 as an example:

1. Split horizon group “access” is created to prevent any communication between the SAP’s part of port 1/1/4 and port 1/1/5 (configured as access port) within the same VPLS.
2. Split horizon group “network” is created to prevent any communication between port 1/1/1 and port 1/1/3 (configured as a network port) within the same VPLS.
3. VPLS 100 is created on 7210 SAS-1 with spoke SDPs on network port 1/1/1 and 1/1/3, and SAPs on 1/1/4 and 1/1/5 as part of this VPLS. CE1, CE2 and CE3 are the customer sites.
4. With this configuration, any communication between ports 1/1/4 and 1/1/5 gets blocked, similarly communication between ports 1/1/1 and 1/1/3 gets blocked but any traffic received on ports (for example, spoke SDPs on these ports) that belong to split horizon group “network” will be switched to ports (for example, SAPs on these ports) that belong to split horizon group “access” and vice versa based on the FDB entries for VPLS 100.

Configuration Guidelines

The following configuration guidelines must be followed to configure a split horizon group.

1. Create a split horizon group in the config prompt. The group name must be unique across the system.

```
7210-SAS1>config#info
#--------------------------------------------------
echo "Split-horizon-group Configuration"
#--------------------------------------------------
    split-horizon-group access create
        description “Block access between access Ports”
    split-horizon-group network create
        description "Block access between network Ports"
    exit
#--------------------------------------------------
7210-SAS1>config#
```

2. Configure ports 1/1/4 and 1/1/5 as access ports and associate these ports with split horizon group “access”.

```
7210-SAS1>config#info
#--------------------------------------------------
echo "Port Configuration"
#--------------------------------------------------
    port 1/1/4
        split-horizon-group access
        ethernet
            mode access
            access
            exit
        exit
        no shutdown
    exit
    port 1/1/5
        split-horizon-group access
        ethernet
            mode access
            access
            exit
        exit
        no shutdown
     exit
#--------------------------------------------------
7210-SAS1>config#
```

3. Configure ports 1/1/1 and 1/1/3 as network ports and associate these ports with split horizon group “network”. The default Ethernet encapsulation for network port is null.

```
7210-SAS1>config# info
#------------------------------------------------
echo "Port Configuration"
#------------------------------------------------
    port 1/1/1
        split-horizon-group network
        ethernet
        exit
        no shutdown
exit
    port 1/1/3
        split-horizon-group network
        ethernet
        exit
        no shutdown
  exit
#------------------------------------------------
7210-SAS1>config#
```

4. Create a VPLS instance 100.

```
#------------------------------------------------
echo "Service Configuration"
#------------------------------------------------
    service
          customer 2 create
          exit
         vpls 100 customer 2 create
            stp
                shutdown
            exit
     sap 1/1/4 create
     exit
     sap 1/1/5 create
     exit
     spoke-sdp 1:1 create
     exit
     spoke-sdp 2:1 create
     exit
        no shutdown
    exit
...
#------------------------------------------------
```

Note: A split horizon on a port must be configured before creating any SAPs associated with that port.

Verification

The following output verifies the split horizon configuration on a 7210 SAS:

```
7210-SAS1# show split-horizon-group
===============================================================================
Port: Split Horizon Group
===============================================================================
Name                              Description
-------------------------------------------------------------------------------
access                            Block access between access Ports
network                           Block access between network Ports

No. of Split Horizon Groups: 2
===============================================================================
7210-SAS1#
```

Execute the below mentioned command to verify the port association with split horizon groups:

```
7210-SAS1# show split-horizon-group access
===============================================================================
Port: Split Horizon Group
===============================================================================
Name                              Description
-------------------------------------------------------------------------------
access                            Block access between access Ports
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Port1/1/4                         10/100/Gig Ethernet SFP
Port1/1/5                         10/100/Gig Ethernet SFP

Ports Associated : 2
===============================================================================
7210-SAS1#


7210-SAS1# show split-horizon-group network
===============================================================================
Port: Split Horizon Group
===============================================================================
Name                              Description
-------------------------------------------------------------------------------
network                           Block access between network Ports
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Port1/1/1                         10/100/Gig Ethernet SFP
Port1/1/3                         10/100/Gig Ethernet SFP

Ports Associated : 2
===============================================================================
7210-SAS1#
```

Appendix: DHCP Management

In This Chapter

This chapter provides information about using DHCP, including theory, supported features and configuration process overview.

The topics in this chapter include:

DHCP Principles

In a Triple Play network, client devices (such as a routed home gateway, a session initiation protocol (SIP) phone or a set-top box) use Dynamic Host Configuration Protocol (DHCP) to dynamically obtain their IP address and other network configuration information. 7210 autoinit procedure also uses DHCP to dynamically obtain the BOF file used for first-time booting of the system (along with IP address required to retrieve the BOF file, the configuration file and the Timos software image from the network). DHCP is defined and shaped by several RFCs and drafts in the IETF DHC working group including the following

- RFC 2131, Dynamic Host Configuration Protocol
- RFC 3046, DHCP Relay Agent Information Option

The DHCP operation is illustrated in Figure 71.

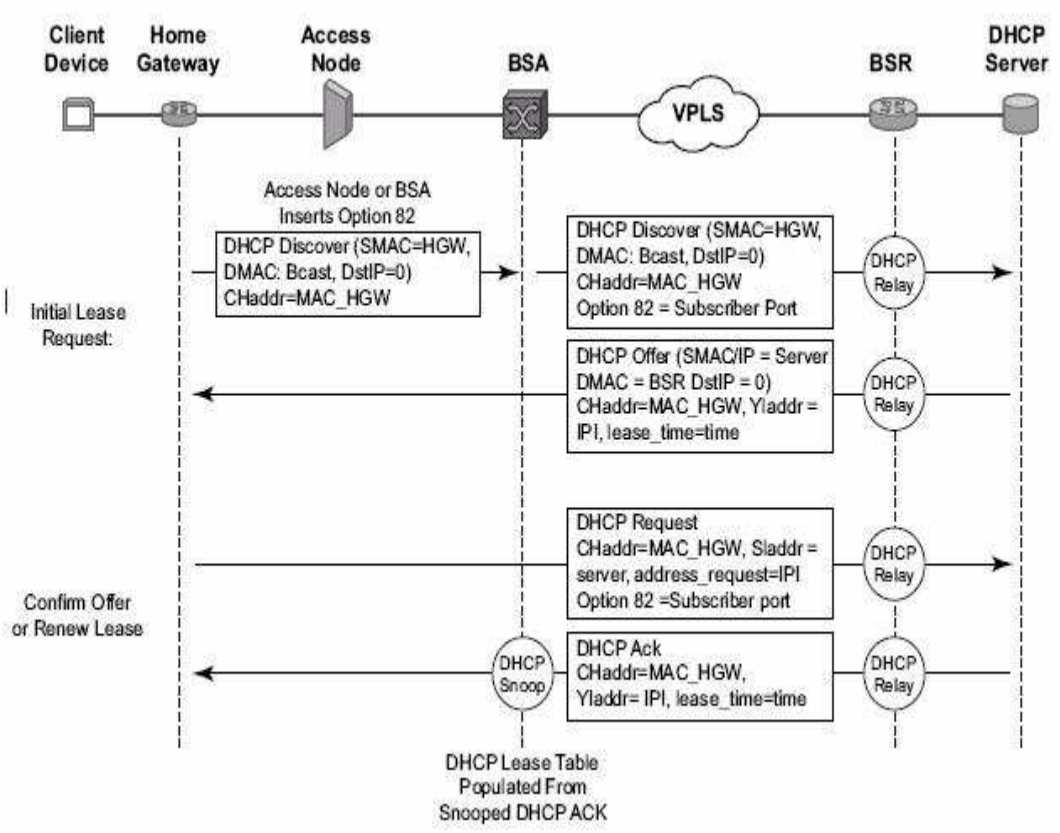

Figure 71: IP Address Assignment with DHCP

1. During boot-up, the client device sends a DHCP discover message to get an IP address from the DHCP Server. The message contains:
   - → Destination MAC address — broadcast
   - → Source MAC address — MAC of client device
   - → Client hardware address — MAC of client device

If this message passes through a DSLAM or other access node (possibly a 7210 SAS device), typically the Relay information option (Option 82) field is added, indicating shelf, slot, port, VPI, VCI and other fields, to identify the subscriber.

DHCP relay is enabled on the first IP interface in the upstream direction. Depending on the scenario, the DSLAM, BSA or the BSR will relay the discover message as a unicast packet towards the configured DHCP server. DHCP relay is configured to insert the giaddr in order to indicate to the DHCP server in which subnet an address should be allocated.

2. The DHCP server will lookup the client MAC address and Option 82 information in its database. If the client is recognized and authorized to access the network, an IP address will be assigned and a DHCP offer message returned. The BSA or BSR will relay this back to the client device.
3. It is possible that the discover reached more than one DHCP server, and thus that more than one offer was returned. The client selects one of the offered IP addresses and confirms it wants to use this in a DHCP request message, sent as unicast to the DHCP server that offered it.
4. The DHCP server confirms that the IP address is still available, updates its database to indicate it is now in use, and replies with a DHCP ACK message back to the client. The ACK also contains the Lease Time of the IP address.

DHCP Features

Using Option 82 Field

Option 82, or the relay information option is specified in RFC 3046, DHCP Relay Agent Information Option, allows the router to append some information to the DHCP request that identifies where the original DHCP request arrives from.

There are two sub-options under Option 82:

- Agent Circuit ID Sub-option (RFC 3046, section 3.1): This sub-option specifies data which must be unique to the box that is relaying the circuit.
- Remote ID Sub-option (RFC 3046 section 3.2): This sub-option identifies the host at the other end of the circuit. This value must be globally unique.

Both sub-options are supported by the Alcatel-Lucent 7210 SAS and can be used separately or together.

Inserting Option 82 information is supported independently of DHCP relay.

When the circuit id sub-option field is inserted by the 7210 SAS, it can take following values:

- *sap-id*: The SAP index (only under a IES or VPRN service)
- *ifindex*: The index of the IP interface (only under a IES or VPRN service)
- *ascii-tuple*: An ASCII-encoded concatenated tuple, consisting of [system-name|serviceid| interface-name] (for VPRN or IES) or [system-name|service-id|sap-id] (for VPLS).
- *vlan-ascii-tuple*: An ASCII-encoded concatenated tuple, consisting of the ascii-tuple followed by Dot1p bits and Dot1q tags.

Note that for VPRN the ifindex is unique only within a VRF. The DHCP relay function automatically prepends the VRF ID to the ifindex before relaying a DHCP Request.

When a DHCP packet is received with Option 82 information already present, the system can do one of three things. The available actions are:

- *Replace* — On ingress the existing information-option is replaced with the information-option parameter configured on the 7210 SAS. On egress (towards the customer) the information-option is stripped (per the RFC).

- *Drop* — The DHCP packet is dropped and a counter is incremented.
- *Keep* — The existing information is kept on the packet and the router does not add any additional information. On egress the information option is not stripped and is sent on to the downstream node.

In accordance with the RFC, the default behavior is to keep the existing information; except if the giaddr of the packet received is identical to a local IP address on the router, then the packet is dropped and an error incremented regardless of the configured action.

The maximum packet size for a DHCP relay packet is 1500 bytes. If adding the Option 82 information would cause the packet to exceed this size, the DHCP relay request will be forwarded without the Option 82 information. This packet size limitation exists to ensure that there will be no fragmentation on the end Ethernet segment where the DHCP server attaches.

In the downstream direction, the inserted Option 82 information should not be passed back towards the client (as per RFC 3046, DHCP Relay Agent Information Option). To enable downstream stripping of the option 82 field, DHCP snooping should be enabled on the SDP or SAP connected to the DHCP server.

Trusted and Untrusted

There is a case where the relay agent could receive a request where the downstream node added Option 82 information without also adding a giaddr (giaddr of 0). In this case the default behavior is for the router to drop the DHCP request. This behavior is in line with the RFC.

The 7210 SAS supports a command trusted, which allows the router to forward the DHCP request even if it receives one with a giaddr of 0 and Option 82 information attached. This could occur with older access equipment. In this case the relay agent would modify the request's giaddr to be equal to the ingress interface. This only makes sense when the action in the information option is keep, and the service is IES or VPRN. In the case where the Option 82 information gets replaced by the relay agent, either through explicit configuration or the VPLS DHCP Relay case, the original Option 82 information is lost, and the reason for enabling the trusted option is lost.

DHCP Snooping

To support DHCP based address assignment in L2 aggregation network, 7210 supports DHCP snooping. 7210 can copy packets designated to the standard UDP port for DHCP (port 67) to its control plane for inspection, this process is called DHCP snooping.

DHCP snooping can be performed in two directions:

1. From the client to the DHCP server (Discover or Request messages) to insert Option 82 information; For these applications, DHCP snooping must be enabled on the SAP towards the subscriber.
2. From the DHCP server (ACK messages), to remove the Option 82 field towards the client. For these applications, DHCP snooping must be enabled on both the SAP towards the network and the SAP towards the subscriber.

Common Configuration Guidelines

The topic in this section are:

Configuration Guidelines for DHCP relay and snooping

The following configuration guidelines must be followed to configure DHCP relay and snooping.

- On 7210 SAS-M and 7210 SAS-X, DHCP snooping is not supported for SDPs
- 7210 SAS devices does not support the ARP populate based on the DHCP lease, assigned to the DHCP client
- 7210 SAS devices does not maintain the DHCP lease assigned to the client
- 7210 SAS devices do not perform IP spoofing checks and MAC spoofing checks based on the DHCP parameters assigned to the client
- MAC learning must be enabled in the VPLS service, for DHCP snooping.
- DHCP snooping is not supported for B-SAPs in B-VPLS services and I-SAPs in I-VPLS services.
- Ingress ACLs cannot be used to drop DHCP control packet.
- DHCP packets received over a SDP cannot be identified and option-82 inserted by the node cannot be removed by the node, in the downstream direction. If this behavior is not needed user should not enable DHCP snooping in the VPLS service, if the DHCP server is reachable over the SDP (either spoke-sdp or mesh-sdp).

Configuring Option 82 Handling

Option 82, or “Relay Information Option” is a field in DHCP messages used to identify the subscriber. The Option 82 field can already be filled in when a DHCP message is received at the router, or it can be empty. If the field is empty, the router shoulMAC learning must be enabled in the VPLS service, for DHCP snooping.d add identifying information (circuit ID, remote ID or both). If the field is not empty, the router can decide to replace it.

The following example displays an example of a partial BSA configuration with Option 82 adding on a VPLS service. Note that snooping must be enabled explicitly on a SAP.

```
*A:7210SAS>config>service#
------------------------------------------------

vpls 2 customer 1 create
```

```
            shutdown
            stp
                shutdown
            exit
sap 1/1/12:100 create
                dhcp                    //Configuration example to add option 82
                    option
                        action replace
                        circuit-id
                        no remote-id
                    exit
                    no shutdown
                exit
            exit
            no shutdown
        exit
----------------------------------------------

*A:7210SAS>config>service#
```

The following example displays an example of a partial BSA configuration to remove the Option 82 on a VPLS service.

```
vpls 2 customer 1 create
            stp
                shutdown
            exit
            sap 1/1/14:100 create      //Configuration example to remove option 82
                dhcp
                    snoop
                    no shutdown
                exit
            exit
```

NOTE: The capabilities available when operating in access-uplink mode/L2 mode and network mode/MPLS mode are different. Correspondingly, not all the standards and protocols listed below are applicable to access-uplink mode and network mode.

Standards Compliance

IEEE 802.1ab-REV/D3 Station and Media Access Control Connectivity Discovery
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 802.1ad Provider Bridges
IEEE 802.1ah Provider Backbone Bridges (Not Available on 7210 SAS-R6)
IEEE 802.1ag Service Layer OAM
IEEE 802.3ah Ethernet in the First Mile
IEEE 802.3 10BaseT
IEEE 802.3ad Link Aggregation
IEEE 802.3ae 10Gbps Ethernet
IEEE 802.3u 100BaseTX
IEEE 802.3z 1000BaseSX/LX
ITU-T Y.1731 OAM functions and mechanisms for Ethernet based networks
draft-ietf-disman-alarm-mib-04.txt
IANA-IFType-MIB
IEEE8023-LAG-MIB
ITU-T G.8032 Ethernet Ring Protection Switching (version 2)

Protocol Support

BGP
RFC 1397 BGP Default Route Advertisement
RFC 1772 Application of BGP in the Internet
RFC 1997 BGP Communities Attribute
RFC 2385 Protection of BGP Sessions via MD5
RFC 2439 BGP Route Flap Dampening
RFC 2547bis BGP/MPLS VPNs
draft-ietf-idr-rfc2858bis-09.txt.
RFC 2918 Route Refresh Capability for BGP-4
RFC 3107 Carrying Label Information in BGP-4
RFC 3392 Capabilities Advertisement with BGP-4
RFC 4271 BGP-4 (previously RFC 1771)
RFC 4360 BGP Extended Communities Attribute
RFC 4364 BGP/MPLS IP Virtual Private Networks (VPNs) (previously RFC 2547bis BGP/MPLS VPNs)
RFC 4760 Multi-protocol Extensions for BGP
RFC 4893 BGP Support for Four-octet AS Number Space
RFC 5291 Outbound Route Filtering
RFC 5668 4-Octet AS Specific BGP Extended Community Capability for BGP-4

CIRCUIT EMULATION (7210 SAS-M Only)
RFC 4553 Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)
RFC 5086 Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)
RFC 5287 Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks

DHCP
RFC 2131 Dynamic Host Configuration Protocol
RFC 3046 DHCP Relay Agent Information Option (Option 82)

DIFFERENTIATED SERVICES
RFC 2474 Definition of the DS Field the IPv4 and IPv6 Headers (Rev)
RFC 2597 Assured Forwarding PHB Group (rev3260)
RFC 2598 An Expedited Forwarding PHB
RFC 2697 A Single Rate Three Color Marker
RFC 2698 A Two Rate Three Color Marker
RFC 4115 A Differentiated Service Two-Rate, Three-Color Marker with Efficient Handling of in-Profile Traffic

IPv6 (7210 SAS-M and 7210 SAS-X only)
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 2464 Transmission of IPv6 Packets over Ethernet Networks
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

RFC 4552 Authentication/Confidentiality for OSPFv3
RFC 5095 Deprecation of Type 0 Routing Headers in IPv6
draft-ietf-isis-ipv6-05
draft-ietf-isis-wg-multi-topology-xx.txt

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 Intermediate System to Intermediate System (ISIS) 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
RFC 3847 Restart Signaling for IS-IS – GR helper
RFC 4205 for Shared Risk Link Group (SRLG) TLV

MPLS - General
RFC 3031 MPLS Architecture
RFC 3032 MPLS Label Stack Encoding
RFC 4379 Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures
RFC 4182 Removing a Restriction on the use of MPLS Explicit NULL

MPLS - LDP
RFC 5036 LDP Specification
RFC 3037 LDP Applicability
RFC 3478 Graceful Restart Mechanism for LDP — GR helper
RFC 5283 LDP extension for Inter-Area LSP
RFC 5443 LDP IGP Synchronization

MPLS - RSVP-TE
RFC 2430 A Provider Architecture DiffServ & TE
RFC 2702 Requirements for Traffic Engineering over MPLS
RFC2747 RSVP Cryptographic Authentication
RFC3097 RSVP Cryptographic Authentication
RFC 3209 Extensions to RSVP for Tunnels
RFC 4090 Fast reroute Extensions to RSVP-TE for LSP Tunnels
RFC 5817 Graceful Shutdown in MPLS and GMPLS Traffic Engineering Networks

MPLS-TP (Transport Profile) [7210 SAS-T and 7210 SAS-R6 in network mode only]
RFC 5586 MPLS Generic Associated Channel
RFC 5921 A Framework for MPLS in Transport Networks
RFC 5960 MPLS Transport Profile Data Plane Architecture
RFC 6370 MPLS-TP Identifiers
RFC 6378 MPLS-TP Linear Protection
RFC 6428 Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile
RFC 6426 MPLS On-Demand Connectivity and Route Tracing
RFC 6478 Pseudowire Status for Static Pseudowires
draft-ietf-mpls-tp-ethernet-addressing-02 MPLS-TP Next-Hop Ethernet Addressing

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) [Only in 7210 SAS-M and 7210 SAS-T access-uplink mode and supported with IPv4 Multicast in network mode]
RFC 2362 Protocol Independent Multicast Sparse Mode (PIMSM)
RFC 4601 Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)
RFC 4604 Using IGMPv3 and MLDv2 for Source-Specific Multicast
RFC 4607 Source-Specific Multicast for IP
RFC 4608 Source-Specific Protocol Independent Multicast in 232/8
draft-ietf-pim-sm-bsr-06.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 1215 A Convention for Defining Traps for use with the SNMP
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 2571 SNMP-FRAMEWORKMIB
RFC 2572 SNMP-MPD-MIB
RFC 2573 SNMP-TARGET-&-NOTIFICATION-MIB
RFC 2574 SNMP-USER-BASEDSMMIB
RFC 2575 SNMP-VIEW-BASEDACM-MIB
RFC 2576 SNMP-COMMUNITY-MIB
RFC 2665 EtherLike-MIB
RFC 2819 RMON-MIB
RFC 2863 IF-MIB
RFC 2864 INVERTED-STACK-MIB
RFC 3014 NOTIFICATION-LOGMIB
RFC 3164 Syslog
RFC 3273 HCRMON-MI
RFC 3411 An Architecture for Describing Simple Network

Management Protocol (SNMP) Management Frameworks
RFC 3412 - Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
RFC 3413 - Simple Network Management Protocol (SNMP) Applications
RFC 3414 - User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)
RFC 3418 - SNMP MIB
draft-ietf-mpls-lsr-mib-06.txt
draft-ietf-mpls-te-mib-04.txt
draft-ietf-mpls-ldp-mib-07.txt

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 3623 Graceful OSPF Restart – GR helper
RFC 3630 Traffic Engineering (TE) Extensions to OSPF Version 2
RFC 2740 OSPF for IPv6 (OSPFv3) draft-ietf-ospf-ospfv3-update-14.txt
RFC 4203 Shared Risk Link Group (SRLG) sub-TLV
RFC 4577 OSPF as the Provider/ Customer Edge Protocol for BGP/ MPLS IP Virtual Private Networks (VPNs)

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)
RFC 4448 Encapsulation Methods for Transport of Ethernet over MPLS Networks (draft-ietf-pwe3-ethernet-encap-11.txt)
RFC 4446 IANA Allocations for PWE3
RFC 4447 Pseudowire Setup and Maintenance Using LDP (draft-ietf-pwe3-control-protocol-17.txt)
RFC 5085, Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires
RFC 5659 An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge
RFC6073, Segmented Pseudowire (draft-ietf-pwe3-segmented-pw-18.txt)
draft-ietf-l2vpn-vpws-iw-oam-02.txt, OAM Procedures for VPWS Interworking
draft-ietf-pwe3-oam-msg-map-14-txt, Pseudowire (PW) OAM Message Mapping
draft-muley-dutta-pwe3-redundancy-bit-03.txt, Pseudowire Preferential Forwarding Status bit definition
draft-pwe3-redundancy-02.txt, Pseudowire (PW) Redundancy

RADIUS
RFC 2865 Remote Authentication Dial In User Service
RFC 2866 RADIUS Accounting

SSH
draft-ietf-secsh-architecture.txt SSH 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

TCP/IP
RFC 768 UDP
RFC 1350 The TFTP Protocol
RFC 791 IP
RFC 792 ICMP
RFC 793 TCP
RFC 826 ARP
RFC 854 Telnet
RFC 1519 CIDR
RFC 1812 Requirements for IPv4 Routers
RFC 2347 TFTP option Extension
RFC 2328 TFTP Blocksize Option
RFC 2349 TFTP Timeout Interval and Transfer Size option
draft-ietf-bfd-mib-00.txt Bidirectional Forwarding Detection Management Information Base
RFC 5880 Bidirectional Forwarding Detection
RFC 5881 BFD IPv4 (Single Hop)
RFC 5883 BFD for Multihop Paths (7210 SAS-M only)

Timing
ITU-T G.781 Telecommunication Standardization Section of ITU, Synchronization layer functions, issued 09/2008
ITU-T G.813 Telecommunication Standardization Section of ITU, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003.
GR-1244-CORE Clocks for the Synchronized Network: Common Generic Criteria, Issue 3,May 2005
ITU-T G.8261 Telecommunication Standardization Section of ITU, Timing and synchronization aspects in packet networks, issued 04/2008. (Not available on 7210 SAS-R6)
ITU-T G.8262 Telecommunication Standardization Section of ITU, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007.
ITU-T G.8264 Telecommunication Standardization Section of ITU, Distribution of timing information through packet networks, issued 10/ 2008.
IEEE Std 1588™-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems. (Not available on 7210 SAS-R6)

VPLS

RFC 4762 Virtual Private LAN Services Using LDP (previously draft-ietf-l2vpn-vpls-ldp-08.txt)

VRRP

RFC 2787 Definitions of Managed Objects for the Virtual Router Redundancy Protocol

RFC 3768 Virtual Router Redundancy Protocol

Proprietary MIBs

ALCATEL-IGMP-SNOOPING-MIB.mib
TIMETRA-CAPABILITY-7210-SAS-M-V5v0.mib (7210 SAS-M Only)
TIMETRA-CAPABILITY-7210-SAS-X-V5v0.mib (7210 SAS-X Only)
TIMETRA-CHASSIS-MIB.mib
TIMETRA-CLEAR-MIB.mib
TIMETRA-DOT3-OAM-MIB.mib
TIMETRA-FILTER-MIB.mib
TIMETRA-GLOBAL-MIB.mib
TIMETRA-IEEE8021-CFM-MIB.mib
TIMETRA-LAG-MIB.mib
TIMETRA-LOG-MIB.mib
TIMETRA-MIRROR-MIB.mib
TIMETRA-NTP-MIB.mib
TIMETRA-OAM-TEST-MIB.mib
TIMETRA-PORT-MIB.mib
TIMETRA-QOS-MIB.mib
TIMETRA-SAS-ALARM-INPUT-MIB.mib
TIMETRA-SAS-FILTER-MIB.mib
TIMETRA-SAS-IEEE8021-CFM-MIB.mib
TIMETRA-SAS-IEEE8021-PAE-MIB.mib
TIMETRA-SAS-GLOBAL-MIB.mib
TIMETRA-SAS-LOG-MIB.mib.mib
TIMETRA-SAS-MIRROR-MIB.mib
TIMETRA-SAS-MPOINT-MGMT-MIB.mib (Only for 7210 SAS-X)
TIMETRA-SAS-PORT-MIB.mib
TIMETRA-SAS-QOS-MIB.mib
TIMETRA-SAS-SDP-MIB.mib
TIMETRA-SAS-SYSTEM-MIB.mib
TIMETRA-SAS-SERV-MIB.mib
TIMETRA-SAS-VRTR-MIB.mib
TIMETRA-SCHEDULER-MIB.mib
TIMETRA-SECURITY-MIB.mib
TIMETRA-SERV-MIB.mib
TIMETRA-SYSTEM-MIB.mib
TIMETRA-TC-MIB.mib
TIMETRA-ISIS-MIB.mib
TIMETRA-ROUTE-POLICY-MIB.mib
TIMETRA-MPLS-MIB.mib
TIMETRA-RSVP-MIB.mib
TIMETRA-LDP-MIB.mib
TIMETRA-VRRP-MIB.mib
TIMETRA-VRTR-MIB.mib

INDEX

T

V